WAVES

NPP & TDM Software Guides

L1
C1
S1
Q10
TrueVerb
PAZ
DeEsser
SuperTap 2-tap
IDR

Mac

serial number

92843NPP3A6163

NPP 3.2

USW379800F092

Read Me First

Read Me First!

Thank you for buying Waves audio products. This page will give you some basic information about the products, where to start, and answer a few common questions.

• In the past, we used hardware copy-protection. In the interest of an easier installation process for you, Waves now uses a challenge/response system. Your product will work full-function for 14 days after installation, giving you plenty of time to authorize your product. This 14 day period is a one-time feature and can not be reset or reinstalled.

• Our Best Recommendation: Follow the Installation instructions on the next few pages, then read the WaveSystem Manual while you play around with one of the plug-ins. Waves processors have extraordinary power in the user interfaces and the WaveSystem offers you tips and hints that will make your work go faster and easier. You wouldn't believe how many people only find out about some of these tips 2 years after they bought a product!

• With almost every product, you can install to either Mac or PC, but only one! As always you can use your product with more than one application on your Mac or PC as long as it uses a supported protocol.

Q: I already own a Wavekey from a previous purchase. Can I still use it?
A: Yes but only with software prior to version 3. Chances are, if you have a WaveKey, you have purchased an upgrade (instructions are on the next page) and you are no longer eligible for tech support for that key.

Q: Can I move the authorization around?
A: Only if you install to an external hard drive. Jaz and Zip are not supported.

Q: What are the system requirements for the plug-ins?
A: See the package itself for requirements. The requirements for the plug-ins vary depending on the host application and plug-in architecture you are using. For the packages which access the processing power of the TDM cards, the performance on ProTools TDM systems is relatively predictable and is determined by the number and type of DSP cards you have. Multitrack native systems (such as Cubase VST) are less predictable because they use the computer's CPU to do all the processing. The CPU not only has to run the plug-ins, it has to run the operating system, the host application, any other applications or system extensions, and other various tasks. Therefore, the faster your system, the more functionality you will gain (i.e. more real-time plug-ins).

Installation
Tech Support

Installation for Macintosh or PC

For everyone, start here

• Inside your package is a **Product Registration** page, which has your all-important **Waves Serial Number**. The page can also serve as a "manual backup" if you are unable to access the Internet. Our online authorization procedure is available 24 hours a day worldwide. If you can't use this method, see the instructions on the Product Registration page.

• Run the installer from the Waves CD supplied in the package. The installer file is named "InstallWaves30" (Mac), or "InstallWaves30.exe" (PC). After installation, your Waves processors will work with full function for 14 days. You have until then to authorize the hard drive of your choice; an external drive would allow you to move your authorizations between machines.

• Have your Waves Serial Number with you at your online computer (does not have to be your workstation),then point your browser to **http://register.waves.com.** Now proceed to either the New Customer section or the Upgrading from WaveKey section.

New Customer (not upgrading from WaveKey system)

• Click "I'm a new customer" to create your account, then select "Register New Product", even if you are upgrading. A new page, "Product Registration" will be displayed. Fill out the page completely, and put your Waves Serial Number in the Product Serial # field.

• Now your account will be created. It may take a little while for your "Manage My Account" page to be updated. Your registered product will be shown on the Registered Product list.

• Proceed to the "Authorization" section on the next page.

Upgrading from WaveKey

• Click "I'm a new customer" to create your account, then select "Register New Product", even if you are upgrading. A new page, "Product Registration" will be displayed. Fill out the page completely, and type in your WaveKey number.

• Now your account will be created. It may take a little while for your "Manage My Account" page to be updated. Your registered product (the older WaveKey package) will be shown on the Registered Product list.

• Now register your upgrade from the Manage Account page by clicking the "Upgrade"button beside your registered WaveKey product; then enter the new Waves Serial Number from inside your package.

Authorization of registered product

• Before you continue with this step, you must have completed one of the two sections on the previous page and have your new product serial number registered, and showing, in your "Manage my Account" page.

• From the Apple Menu (Mac) or your Start menu (PC), select the Waves Authorizer (or locate it in the Waves folder where you installed the Waves software).

• Carefully select the hard drive you wish to authorize; if you wish to have portability between workstations, install to an external hard drive (removables such as Jaz or Zip are not supported).

• Copy the challenge by clicking the "Copy Challenge" button. Clicking the "Web Page" button will take you to "register.waves.com" if you have an internet connection already established. If you are online, just leave the Waves Authorizer running so you can paste in the response in a few moments.

• Access your Waves account which you created by following instructions on the previous page.

• Click the "Authorize" button beside the Registered Product line (which will show your serial number). Paste or type the challenge into the Challenge field, and click "Submit Authorization".

• The system will reply with your Response, which you can select and copy, then paste back into the Authorizer. The "OK" button will become active to show the Response was correct. Click the OK button to complete the authorization.

<u>After installation is complete:</u>

• Increase the memory to your applications as indicated below in the Memory Allocation section.

• Launch the application you wish to use.

• (Note for Pro Tools only): During launch, Pro Tools will display a dialog with several options: [1] Don't use AudioSuite; [2] Extended backward compatibility; [3] Don't ask me again. If checked, the first option no Waves processors will be in the AudioSuite menu (for speed and menu length advantage). The Extended compatibility option will load any old sessions which used the "+" plug-ins (such as +Q10); however, you will see those plug-ins listed <u>twice</u> in the menu,so don't check this unless you need it. The final option prevents this dialog box from being shown each time you launch Pro Tools. You can force the dialog to be shown by simply holding down the space bar during any launch of Pro Tools.

• <u>In all applications</u>, you will be prompted to locate the folder where the Waves plug-ins are found. Navigate to the Waves folder, then Choose the sub-folder "Waves Plug-ins". This creates a preference file (such as WaveShell-DAE Preferences, or WaveShell-MAS Preferences) which "points" to the Waves Plug-ins folder.

Memory Allocation

For each application installed, open the application's information dialog box and add the following memory:

• ProTools and all other DAE applications - Add at least 8MB to DAE (located in the System Folder in the DAE folder) to the minimum and preferred size. We recommend adding at least 16MB. For more information, check the recommendations of your application. For Pro Tools|24 Mix systems, make sure that DAE has at least 10MB "headroom" on large sessions.

• Cubase VST - Add 750KB per plug-in instance to VST.

• All other native applications - Add 3MB to the application's minimum and preferred size.

Technical Support

In order to provide you with the best technical support possible, Waves has a variety of technical support services available.

You must be registered to be eligible for technical support, and must be prepared to give your serial number, User ID, or other requested information in order to proceed with technical support services.

Technical Support options:

•First, check the read me files in the Waves folder and the rest of this manual. The answer could be right under your right (or left) thumb.
•Go to our website at www.waves.com for updates and FAQs.
•Email our tech support team directly at the address listed below.
•If you live outside of the US, you can contact your local Waves dealer. Most local dealers can provide you with authorizations, upgrades and technical information. A list of worldwide dealers and distributors is available on our website.

We invite you to check out www.waves.com for technical information and all the latest information about Waves products.

Worldwide technical support

(Monday-Friday, 9am-6pm east coast)

tech support:
support@waves.com
306 West Depot, Suite 100
Knoxville, TN 37917 USA
tel: 865-546-6115
fax: 865-546-8445

For sales and pre-sales product information:
sales@waves.com

WaveSystem

Table of Contents

Chapter1 ..... Quick WaveSystem overview 2

Chapter 2 ..... Register now to get updates 3

Waves controls 4
Button-type controls 5
Graphic-type controls 6
Selection of controls 8
TAB functions 9

Chapter 3 ..... Using Setups - Macintosh ... 10

Load and Save 10
Loading Setups 11
Saving Setups 12
Deleting Setups 12
A/B compare, Copy to A (or B) 13
Editing and Pasting 13

Chapter 4 ..... Using Setups - PC ... 15

Load and Save 15
Loading Setups 16
Saving and Deleting Setups 17
A/B Compare, Copy nto A (or B) 18

Chapter 5 ..... Meters, Input, and Output ... 19

Peak-hold and Clip lights/Counters 19
Phase reverse and channel swap 19
Trim/Clip 20
Bypass and Bypass gain 20

Chapter 6 ..... IDR (Increased Digital Resolution) ... 21

Chapter 7 ..... Component plug-ins and the WaveShells ... 22

Chapter 1 - Quick WaveSystem overview

This collection of manuals contains information for everything in your bundle. There may be additional manuals for some plug-ins not in your package. Use them at your convenience when trying them as demos.

What's in this manual

- How-to-use info on all the controls that are common to all plug-ins (plus special tips for power users and more).
- Info on how component plug-ins and WaveShells work.
- WaveKey Utility instructions.

How things work (Apple Macintosh)

After a manual install, or a Waves Installer procedure, here's where things should be, and how they work. See your WaveShell Read Me file for complete details and last-minute notes:

All Waves plug-ins are in a folder named "Waves plug-ins", which is inside the Waves folder, located on your internal startup disk. They don't go in your application's plug-in folder.

Instead, **A WaveShell** is installed into your application's plug-in folder. The WaveShell connects your application to the Waves plug-ins as a software gateway. For example, Premiere has WaveShell-P in the "Adobe Premiere™ Plug-Ins" folder; VST has the WaveShell-VST in the "VstPlugIns" folder. You may have several WaveShells installed.

The WaveKey is attached to your computer; it is required for all products except AudioTrack (native) and WaveConvert.

It works like this: your program uses its WaveShell to load the plug-ins. Each WaveShell "points" to the Waves Plug-Ins folder. It is that simple to use, and lets you "buy once, use many places".

How things work (Windows for PC)

After a manual install, or a Waves Installer procedure, here's where things should be, and how they work. See the Waves Readme located in Start Menu -> Programs -> Waves for more information.

All Waves plug-ins are in a folder named "plug-ins", which is inside the Waves folder, located in the Program Files folder. They don't go in your application's plug-in folder.

It works like this: The host applications (e.g. Sound Forge) access the plug-ins using the DirectX API. The plug-ins are registered in this format, therefore any application that supports the DirectX audio plug-in format will have access to them.

The WaveKey is attached to your computer; it is required for all products except AudioTrack (native) and WaveConvert. The Wavekey also requires a driver that should be installed automatically.

If you have installation problems be sure to check our website for new versions of the Plug-Ins, the DirectX extensions and the Wavekey Driver.

Chapter 2 - Updates & Demos

Check our website —**www.waves.com** — for the latest versions of all our software. You just download and use any update. The web page also has demos of Waves applications and plug-ins, Frequently-Asked Questions and answers, User's Guides for all products in Adobe™ Acrobat™ (.pdf) format, and much more.

What is a software plug-in?

A plug-in is a software add-on that gives enhanced capabilities to the host software with which it is being used. Think of it as a software accessory. Waves plug-ins do not function as stand-alone applications, they must always be accessed from within the host software, and in the case of digital audio, specialized audio hardware (usually associated with the host system) is required. Once the plug-in has been installed, it may be called up directly via the appropriate menu in the host application. See your host application manual for information on how to access plug-ins. **See the Read Me file in the package or on your disk for Waves plug-in installation instructions.**

Waves controls

This will save you hours! Read while using a plug-in. The controls described here are common to all Waves plug-ins and applications. So, you only have to learn the WaveSystem one time!

WaveSystem controls are of two main types, **Button-type** and **Graphic-type.** Please note your application may not support all these functions (for instance, a sequencer may not allow any keyboard controls). See the "Notes" chapter and readme for the latest info about your application.

Graphic-types can be cross-hairs, small triangle markers, or standard faders. **Button-types** are always boxes with text or numbers in them.

All controls can be in one of three conditions:

1— **Not selected:** the control is not the target of any user entry.
2— **Selected:** the control is the target of mouse control entry only and not keyboard entry. (If the control is a button, it will be depressed, and a marker or a slider will be highlighted).
3— **Selected and Active:** the control is the target for both mouse and keyboard entry. (A yello outline frame will be added around the Active control and blue around Selected).

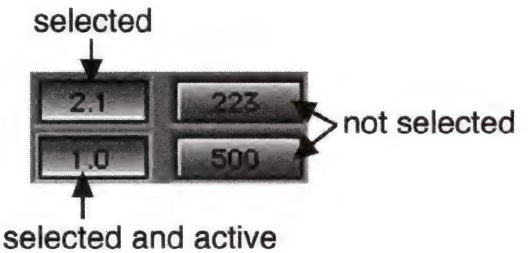

Normally, a control is either 'Not Selected' or 'Selected and Active'. To make a control 'Selected and Active', you simply click it with the mouse, or navigate to it using TAB. This will deactivate the previously Active control.

When multiple controls are selected (using shift-click, or click-drag-select), all of them will be 'Selected' but *only one of them will also be Active* (normally the first one). Within a group of selected controls, you can Activate a certain control by clicking on it, or use the TAB key to navigate to it.

Button-type controls

Buttons can be toggle switches, or allow direct numerical control (and then they are called **Value Windows**).

Examples of Button switches:

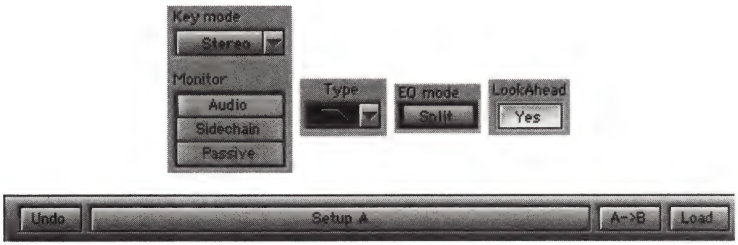

These rectangular regions allow switching between two or more modes. To change mode of a button, simply click on it with the mouse. The current setting of the button will be indicated on the button. In some cases certain modes (e.g. bypass or special monitor modes) will illuminate the button to warn that you may not hear what you might expect according to the settings of other controls.
Value Window Buttons. These rectangular areas look identical to buttons except that they display the current numerical value of the control.

Examples of Value Window buttons:

Click once on the button you wish to edit. To change the value, use the mouse, the arrow keys, or the number pad.

Using the mouse, click-and-drag on the number to change it. Some move up/down, some move left/right. This allows adjustment of more than one control at the same time, such as Gain and Frequency in the Q10. Small arrows will appear in the button to show which way to move the mouse, shown here:

Using the arrow keys, press up/down or left/right to adjust the numbers in the smallest step. Hold the arrow keys down for faster changes.

Using the number pad, enter the number and press Enter to finish, or Escape to leave the original value unchanged. You may also press TAB, which will enter your value and move to the next button. If you enter an out-of-range number, the system beep will sound and the value remains highlighted. You may enter a valid number or press Escape.

Linked Value Windows. In some processors, these unmarked square buttons fill the space between certain pairs of value windows, and allow both to be adjusted with a single click-and-drag. Shown here are Link buttons from the C1comp/gate; the left illustrating control of the Comp Threshold and the GateOpen at the same time:

Linked buttons. In a similar way, adjacent buttons may be linked by an unmarked "link button". Clicking on this link button will cause adjacent buttons to toggle through their modes at the same time (such as EQmode in C1).

(Please note there are no linked buttons or linked value windows in the PC version of the C1-Compressor/Gate.)

Graphic-type controls

Faders. You may click on the fader itself or anywhere within the fader "track". The numerical value of slider control settings is displayed in a window above the slider path. The fader will remain selected; to deselect, click in an open space.

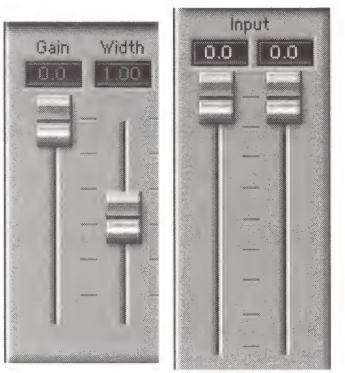

Pairs of controls are provided in some processors, (shown on the right from Q10) for controlling gains of individual stereo channels. In this case, they may be linked by clicking-and-dragging between controls — this will move both up and down together. They may be offset separately, and when dragged together will keep the difference in dB.

Grab Markers. These provide an alternative mode of operation for certain controls. These triangular markers near meters or display graphs, or the cross-hair markers on graphs, may be "grabbed" by clicking and dragging. Generally, triangular markers move along a horizontal or vertical line. In some cases, they may be grabbed by clicking anywhere within their line of movement, not just on the marker itself.

Examples of Grab Markers (triangles and cross-hairs):

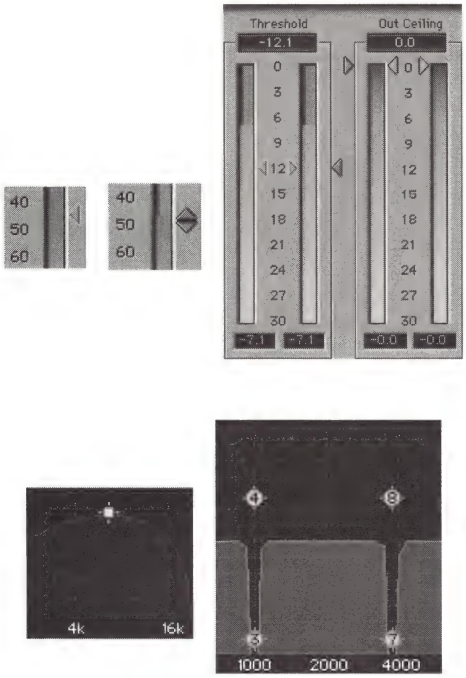

Cross-hair markers on graph displays can usually be moved in two dimensions, both horizontally and vertically. In the horizontal direction, they often control the parameter along the bottom of the graph (e.g. frequency), although other options (e.g. Q for the Q10) may sometimes be invoked by pressing a modifier key (usually the Option key). In the vertical direction they control a second parameter; for example, 'Gain' in the Q10, or the 'Q' in the C1's sidechain EQ.

Selection of controls

One of the most powerful features of Waves audio processors is the ability to select and adjust many controls at the same time. You may select any number of controls for adjustment, copying or pasting together in one operation.

A simple but powerful example is to move many bands of EQ in the Q10 at the same time, either in the graph, or by selecting multiple Frequency buttons.
Drag-select the desired buttons, Value windows and/or grab markers by clicking and holding the mouse at a point outside the desired controls, and dragging the mouse to form a rectangle to include the desired controls. In the following illustration, we're selecting the entire Comp/Exp section of the C1 for pasting into Setup B:

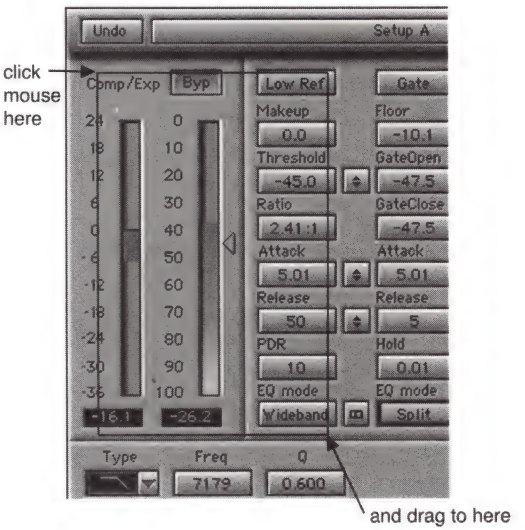

Letting go of the mouse button, you will have selected all controls intersected by the rectangle (and the buttons will appear as 'Selected' by showing a different border).

After selection, clicking or dragging on any single one of the selected controls will have the effect of clicking or dragging simultaneously on all the selected buttons. **Be warned** that the scaling, ranges and calibration of different value windows may be different, so that the resulting linking of controls of different types may not always be what you expect.

For example: suppose you have loaded the "AM Radio" setup in Q10, and want to change the hi-frequency cutoff, which is made from 5 bands of EQ. To do this, you need to change the Frequency of 5 bands at the same time. Simply click-and-drag to select the buttons (shown here).

Then click and drag left/right (horizontally) on any of the selected controls to move all 5 bands at the same time.

In general, it is best to use clicking on several selected buttons only if the buttons are of similar type (such as the in/out buttons or Freq controls of several Q10 bands).

TAB functions

TAB will move the 'selected' status to the next control or Value Window,
shift-TAB goes in reverse order,

In addition the Macintosh version has the following Tab functions.

option-TAB moves "down", where possible,
shift-option-TAB moves "up", where possible.

If you have several Value Windows selected (either thru shift-click or click-drag-select), you can TAB and shift-TAB to go forward and backward through the selected controls. Only one of these controls can be Active, and only the Active control will accept direct numeric entry. This is the Value Window whose value will be changed by the arrow keys, numeric entry, or the mouse. To move to the next Selected control, just press the TAB key.

Load and Save

Waves' Load/Save functions have great flexibility, configurable menus, and fast access in sessions.

To use the Load or Save button, click and hold to see the popup menu.

Here are the main features of the menus:

Factory Presets These are permanent presets in the Load popup menu, created by Waves. They can not be changed or deleted. Each component plug-in may have different Factory Presets.

User Presets You can make the current settings of the plug-in become a preset in your Load menu. Users presets can be changed and deleted.

Setup Files Setup files may hold more than one setup. For example, a single file can contain all the setups for a session. When you open a Setup file, all setups inside it become part of your Load popup menu for fast access.

For example, if you load the "Multimedia tools" file for the Q10, all the setups in the file will be added to your Load menu.
Another very good example, useful for multi-setup files in multi-plug-in sessions, is to open the same Setup file on all instances of the plug-in that is in your session. Save all settings you create as setups into that file. This makes all the setups in that file available very quickly in every insert of that plug-in.

Loading Setups

Click and hold on the Load button to see the Load popup menu. Three options may be available; if a choice is not appropriate or available, it will be gray.

The **first** option is the "Open Setup file...". Select it to open any Setup file, whether from the Library or your own.

The **second** option is the currently open Setup file and its setup(s).

The **third** option includes the Factory Presets.

The **fourth** option includes the User Presets.

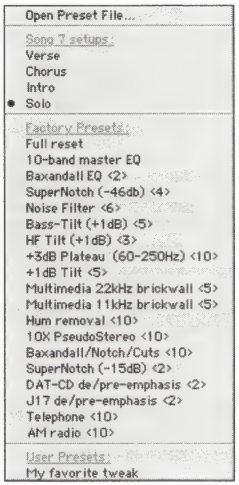

Example:

- From the Q10 plug-in, click and hold the Load button, then select the 'Open Setup File...' option; navigate into the Q10 Setup Library.
- Double-click the 'EQ filters' library file to open it, or highlight it and click the Open button.
- The first setup in the file will be loaded. The first setup in any file you open will be loaded automatically (whether the file has one or many setups). The files are backward-compatible and can be opened with previous versions of Waves plug-ins.
- Click and hold the Load button again.
- In the middle of the menu you will see all the setups that are contained in the 'EQ filters' file.

Saving Setups

Click and hold on the Save button. Four options may be available; if a choice is not appropriate or available, it will be gray.

The **first** menu item will lets you start a new Setup file. There are two prompts, first for the Setup filename, then for the setup name. If you are going to save only one setup in each Setup file, you don't have to name the setup; click OK to complete the save. You might create a folder to hold several setup files for a project.

The second menu item will overwrite the current settings into the loaded setup. This item has the name of the Setup file, followed by the name of the setup itself, all on one line.

Third on the menu will save the current settings as a new setup into the Setup file that is open. (If one is not open, the option is gray.) You will be prompted to give this setup a name.

The **last** menu item will save the current settings into a User Preset that will always be in your Load menu, until you delete it. You will be prompted to give this User Preset a name. Your User Presets are stored in the plug-in's preference file. For instance, if using the Q3 in mono (3 bands of Q10), the User Presets are in the file "Qn (3M) Preferences", in the *System : Preferences : Waves* folder.

Example:

- We will assume you still have the "EQ filters" file open. Select one of the setups.
- Make some changes to the settings, if you wish, then click and hold the Save button.
- The first menu item will allow you to save your tweeked preset to a new setup file.
- The second menu item will save a new setup to the setup that is open.
- The third menu will allow you to save the new tweeked settings to the current setup file. You will be prompted for a new setup name.
- The fourth menu item will save the current preset into the load menu. Here it can be accessed at any time without opening a separate file.

Deleting Setups

You can delete a User Preset, or a setup that is within a Setup file. Factory Presets and Setup Library files can not be deleted or written to.

- Hold the command key down (the "Apple" key).
- Click and hold the Load button to see the popup.
- While still holding the command key, select the Preset or Setup to delete.
- A confirmation box will appear to allow you to cancel or "OK" the operation.
- *Some applications will have different methods for deletion; check the Notes and Read Me files for your application.*

A/B compare, Copy to A (or B)

The **Setup A/Setup B** button may be clicked to compare two settings. If you load a setup in the **Setup B** position, this will not affect the setup loaded into the **Setup A** position, or vice-versa.

If you want to slightly modify the settings in Setup A, you can copy them to Setup B by clicking on the **Copy to B** button, then alter Setup A and compare it with the original in Setup B. When you are listening to Setup B, you can Copy it back into Setup A.

The name of the current setup will be shown in the title bar, (if supported by a platform), and will switch as you change from Setup A to Setup B.

Note: An asterisk is added to the preset name when a change is made to the preset.

Editing and Pasting

The simplest way to develop new customized setups for your own use is to modify the values of controls from a single setup in the Factory Presets, a Setups Library, or from setups you may have saved.

However, there is another way of compiling new setups from existing ones, by editing with copy/paste. Editing allows you to take "bits and pieces" from different setups and incorporate then all into a single new setup. You can combine some or all of the button, Value Window and grab marker settings from two or more setups of a plug-in.

Editing is most useful when different control areas are set to do completely independent tasks. For instance, if 3 bands of the Q10 are being used for a High Tilt, but you need a steep low frequency cut in the bass, you could copy the 3-band High Tilt and paste it into 3 unused bands in the the low cut setup.

You can do editing by compiling the setup you want in Setup A and loading the setups you wish to edit from, one at a time, into Setup B. To paste the settings of a number of control buttons, value widows and/or grab markers from Setup B to Setup A, proceed as in the following example.

Editing example: Suppose you have the 'AM radio' setting loaded into Q10's Setup A. In Setup B, you've loaded the 'High Tilt +1dB' (both are in the Factory Presets, so you can try this easily). We'll combine the low cut of the AM radio with the High Tilt.
First, in Setup B, select the first 3 bands —which make the High Tilt— by selecting the in/out buttons and all Value windows for those 3 bands. Simply click and hold the mouse button at a point outside the desired controls and drag the mouse to form a rectangle to include the desired controls. (In this example, you could also drag-select the Grab markers in Q10).

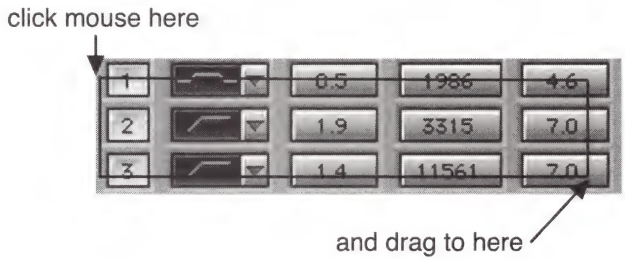

Then press the "c" key on the keyboard (not command-C) to copy the control settings.

Next, switch to Setup A, select three bands (bands 8, 9, and 10) in the same way, then press the "v" key on the keyboard (not command-V) to paste the settings.

Please note the command key is reserved by most applications, so it is not used.

Finally, turn off bands 6 and 7 because they are not used. You'll have bands 1-5 creating a steep low-cut filter, and bands 8-10 providing the High tilt. You still have bands 6 and 7 for more EQ tweaking, or you could use a Q8 to save a little processing power.
Save the newly edited setup as a User Preset, or as a new Setup file.

Special note for any processors with split stereo adjustments, such as Q10): you can copy one channel's settings into the other by the same method. This is done by selecting one stereo channel, say left, selecting controls in a given area, press c on the keyboard to copy, select the other stereo channel (e.g. right), select the same controls again, and press v on the keyboard to paste.

Load and Save

Waves' Load/Save functions have great flexibility, configurable menus, and fast access in sessions. In fact, while Previewing you can load setups while the audio is playing for very easy comparisons.

To use the Load or Save button, click and hold to see the popup menu.

Here are the main features of the menus:

Factory Presets These are permanent presets in the Load popup menu, created by Waves. They can not be changed or deleted. Each component plug-in may have different Factory Presets.

Setup Files Setup files may hold more than one setup. For example, a single file can contain all the setups for a session. When you open a Setup file, all setups inside it become part of your Load popup menu for fast access.

For example, if you load the "Multimedia tools" file for the Q10, all the setups in the file will be added to your Load menu.
Another very good example, useful for multi-setup files in multi-plug-in sessions, is to open the same Setup file on all instances of the plug-in that is in your session. Save all settings you create as setups into that file. This makes all the setups in that file available very quickly in every insert of that plug-in.

Loading Setups

Click and hold on the Load button to see the Load popup menu. Three options may be available; if a choice is not appropriate or available, it will be gray.

The **first** option is the "Open Setup file...". Select it to open any Setup file, whether from the Library or your own.

The **second** option is the currently open Setup file and its setup(s).

The **third** option in the right section of the menu includes the Factory Presets.

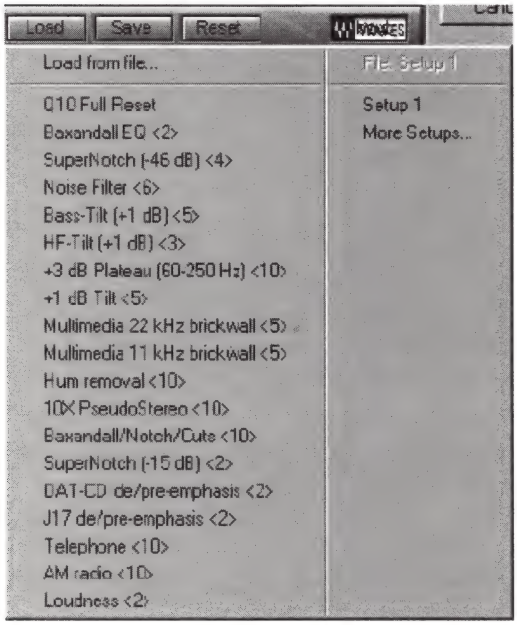

Example:

- From the Q10 plug-in, click and hold the Load button, then select the 'Load from file...' option; navigate into the Q10 Setup Library (located in Program Files -> Waves -> Plug-Ins -> Setup Libraries).
- Double-click the 'EQ filters' library file to open it, or highlight it and click the Open button.
- Click and hold the Load button again.
- In the middle of the menu you will see all the setups that are contained in the 'EQ filters' file.

Saving Setups

Click and hold on the Save button. Three options may be available; if a choice is not appropriate or available, it will be gray.

The **first** menu item will lets you start a new Setup file. There are two prompts, first for the Setup filename, then for the setup name. You must provide a name for both the setup file and the preset. Click OK to complete the save. You might create a folder to hold several setup files for a project.

Second on the menu will save the current settings as a new setup into the Setup file that is open. (If one is not open, the option is gray.) You will be prompted to give this setup a name.

The **third** menu item will overwrite the current settings into the loaded setup. This item has the name of the Setup file, followed by the name of the setup itself, all on one line.

Example:

- We will assume you still have the "EQ filters" file open. Select one of the setups.
- Make some changes to the settings, if you wish, then click and hold the Save button.
- The first menu item will allow you to save your tweeked preset to a new setup file.
- The second menu item will save a new preset to the EQ Filters file. You will be prompted for a preset name.
- The third menu item will save your new settings to the currently open preset.

Deleting Setups

You can delete individual presets within a Setup file.

- Click on the Load button and select "Load for file...:
- Highlight the setup file that contains the preset you wilsh to delete.
- A list of the presets contained in the setup file will appear in the box titled "Preset Name".
- Select the preset you wish to delete and click "Delete".
- You can also delete a preset from the currently open setup fil by selecting "More setups" from the load menu.

A/B compare, Copy to A (or B)

The **Setup A/Setup B** button may be clicked to compare two settings. If you load a setup in the **Setup B** position, this will not affect the setup loaded into the **Setup A** position, or vice-versa.

If you want to slightly modify the settings in Setup A, you can copy them to Setup B by clicking on the **Copy to B** button, then alter Setup A and compare it with the original in Setup B. When you are listening to Setup B, you can Copy it back into Setup A.

The name of the current setup will be shown in the title bar, (if supported by a platform), and will switch as you change from Setup A to Setup B.

Note: An asterisk is added to the preset name when a change is made to the preset.

Chapter 5 - Meters, Input, and Output

Each Waves processor will not have all the features described in this section. For purposes of this example, we use the Q10 input/output section because it has all the features.

Some processors will have input and output faders, some have only a single Master output fader. All have peak-hold meters. In the illustration shown here, you can see all of the features and how to use them.

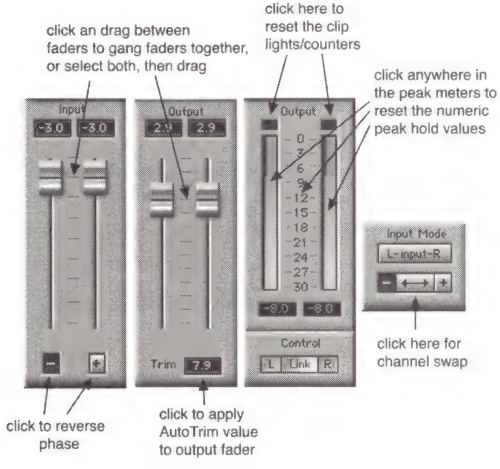

Peak-hold and Clip lights/Counters

Below the peak meters are numbers that show the highest peak in dB since the last reset. To reset these peak hold values, just click directly on the meters.

On most processors, the Clip lights will turn red if any single sample reaches 0dBFS (full scale). A few processors (+S1, you should read the L1 User's Guide for details on setting its output level.

Phase reverse and channel swap

Phase reverse and channel swapping takes place before any other controls in a processor, including gain.

Usually located below the Input Faders are three small switches. The two on the left and right edges control the relative phase of each channel. In the default position, the button is gray, and the plus sign (+) indicates that the phase for that channel is unaltered. To reverse (invert) the phase on a channel, click the phase button. It will turn red and show a negative sign (-) to indicate the phase has been reversed.

Between the phase buttons is the Channel swap button. In the default position, the button is gray and a double-headed arrow indicates that the Left Input feeds Left Output, and the Right Input feeds Right Output. To reverse the channels, click the button. It will turn red and show two long arrows, each pointing from one channel to the other to show the inputs have been swapped.

Trim/Clip

On some processors there are two small displays (in dB) instead of separate Clip lights and Autotrim functions. As in the clip lights, the **Clip** display will turn red and also show the value in dB by how much the gain must be reduced to avoid a clip, as seen in the first graphic.

To adjust gain, just click on the **Trim** window and it will automatically adjust the output fader by the value shown in the Clip display. You can also drag up and down on the Trim display to manually set the gain.

The clip indicator is always active, including during file-based Processing (whether destructive or constructive). This is a great advantage in determining whether a clip occurred during a file-based Process.

Bypass and Bypass gain

Some platforms don't support Bypass. You may use the Setup A/B button to switch to a "flat" Setup B to serve as a bypass.

When you first click the Bypass button, the plug-in will be bypassed as if it is not being used. This allows you to compare your Setup to the "straight" sound. Click the Bypass button again to put the processor in-line.

At the top of some processors near the Bypass button is a Bypass Gain window, labeled BP, that allows you to set the listening volume when the processor is in Bypass mode.

For instance, this is very useful when comparing the quality of a processed sound to the original and you wish to keep the average level about the same.

Chapter 6 - IDR

IDR is (Increased Digital Resolution)

Waves IDR is a great advance in the technology that helps preserve the highest quality during processing. IDR is a noise-shaped dithering system that can be used in consecutive processes, and unlike some systems, is completely compatible with any other previously-dithered files, such as some DAT machines might create using Sony SBM, or, from Apogee A/D converters using UV22.

To use, simply set the IDR button to the bit depth of your on-disk file, or for the output desired. For almost all users, this is 16bit (CD and DAT). In other words, if you are using a platform that stores all the audio files in 16bit, then set IDR for 16bit. In higher-resolution platforms you use dithering only on a Master Fader, as the last plug-in (and generally part of +L1). For multi-channel discreet-output platforms, you can use the IDR plug-in on each output, set to the desired bitdepth. For further information, please read the Notes section and Read Me files for the application you are using.

The IDR function, whether part of a Master plug-in, or used as the small separate IDR plug-in, uses **type1/ Normal** settings. Only the L1 offers more IDR options. For in-depth technical information, consult the L1 User's Guide chapters about IDR.

Chapter 7 - Component plugs and the WaveShells

R-Compressor [stereo]
R-Compressor [mono]
Renaissance EQ2 [stereo]
Renaissance EQ2 [mono]
Renaissance EQ4 [stereo]
Renaissance EQ4 [mono]
Renaissance EQ6 [stereo]
Renaissance EQ6 [mono]
S1 Shuffler [stereo]
S1 Imager [stereo]

Waves makes the most of your hardware by using only what you need. This is true with any system, whether with DSP chips or the computer's CPU (for host-based software, also called 'native').

WaveShell technology lets us 'split' the plug-ins into smaller components. For example, the Q10 lets you choose from 1 to 10 bands of EQ. Most importantly, this saves your processing power, and in many cases, speeds up the processing time too.

Nearly all plug-ins break into components. You'll see them in your menu, as shown to the left.

To use the components: select just the processing that you need for a job. If you find that you need the 'next size up', like more bands of EQ, just save your settings in a Setup file, then reopen it in a 'larger' component.

Not all plug-ins break into every possible configuration; for instance, there is no +L1 (mono), only stereo. However, as new component options are added, they will be noted in the Read Me file with any new plug-in updates. You can always get the latest versions from our web page **www.waves.com** or from other forums. See the Registration page at the beginning of this manual for more addresses.

L1
UltraMaximizer

Table of Contents

Chapter 1 ..... About the L1 3

L1 component plug-ins 3

Chapter 2 ..... Some secrets of digital audio 4

About maximum level 4

About maximum resolution 4

Why use them, and what are they? 5

A bit about IDR (Increased Digital Resolution) 5

Chapter 3 ..... Using the L1 6

The peak limiter section 6

The IDR section 7

Basic control of IDR 8

An example of IDR resolution improvement 9

Type1 9

Type2 10

Noise-shaping options 10

Chapter 4 ..... Important IDR Info 11

Recommended IDR settings 12

Chapter 5 ..... 16-bit (and higher) mastering (CD/DAT) 13

Chapter 6 ..... 8-bit multimedia mastering 14

First: ..... do all other processing, including sample rate conversion 14

Second: ..... choose IDR settings 14

For lowest noise 14

For higher resolution 14

For highest resolution 14

Third: ..... reduce the bitdepth 14

Chapter 7 ..... Use the L1 last 15

Chapter 8 ..... Important notes on digital clipping 16

Peak-normalized signals 16

Practical Precautions 17

Chapter 9 ..... L1 Controls descriptions 18

Quantize 18

16-bit 18

8- and 12-bit 18

18-, 20-, and 22- bit 18

24-bit 18

Dither 18

Shaping 18

Digital/Analog Domain 19

Chapter 10 ..... Recommended settings 20

Chapter 1 - About the L1

The L1-Ultramaximizer is a highly sophisticated audio processing tool kit combining an advanced peak limiter, a level maximizer and a high performance re-quantizer based on the implementation of Michael Gerzon's IDR (Increased Digital Resolution) noise-shaped re-dithering process. Why do we call it the Ultramaximizer? Because it maximizes both the level of the digital signal and the resolution of the final file, working on both 'ends' of the digital word output.

New in the L1 is 48-bit (double precision) processing, offering improvements for all resolutions, and the opportunity to dither to a 24-bit output for archives and DVD. The L1 offers superb requantization for all bitdepths, including 24, 20 , 16, 12, and 8-bit outputs. The development of L1's lookahead peak limiter provides the mastering engineer with the capability to increase soundfile resolution and production master levels with precise control and dithering options. The algorithms are carried out to their fullest potential in the L2 hardware from Waves, although they have their roots in the L1.

Waves IDR™ process brings more choices, greater control, and unmatched compatibility to the mastering environment, whether for high-resolution CD or low-resolution multimedia. IDR offers two dither types and three noise-shaping curves for optimal processing in a wide variety of applications and with a wide choice of source material.

While the operation of conventional limiters is well understood, the limiter section of the L1-Ultramaximizer is capable of a very fast, overshoot-free response, and once the limiter threshold has been set, the user can then go on to define the actual peak level that the processed signal will reach. Once set, limiting and level rescaling becomes a one-shot process. For mastering purposes, the peak level of the processed signal would normally be set to 0dB, or just below 0dB.

Because a typical digital audio file of music contains many high intensity, short duration peaks, simple normalization of the file may still result in a low average signal level. Using the L1-Ultramaximizer however, it is generally possible to significantly increase the average signal level of a typical audio file without introducing any audible side effects. In the event that a deliberately limited sound is required however, there is more than adequate range of the limiter parameters to recreate 'vintage' effects such as level pumping or severely limited dynamic range.

The Ultramaximizer is specifically designed for mastering, digital editing, multimedia, and any application that requires limiting and/or requantization of the digital signal with highest quality. In order to ensure the maximum possible resolution of a processed signal, it is very highly recommended that the L1 be used last in the processing chain. Failure to observe this will not prevent L1 from working, but you should be aware that both the absolute brickwall limiting AND the benefits of IDR requantization will be compromised and will need re-limiting to maintain the original level.

L1 component plug-ins

On most host applications and platforms, you can select just the audio processing you need and use only the power necessary to do the job. The L1 has several plug-in components in the menu.

As of this printing, here are the current components, plus a few examples for each.

- **L1 limiter** — mono/stereo wideband limiter without IDR; *(track insert, voice or instrument limiter, SFX maximizer)*
- **+L1-Ultramaximizer** — the "full" mastering plug-in with limiter and all IDR options. *(CD mastering, multimedia authoring)*

Chapter 2 - Some secrets of digital audio

In order to make the best use of the L1-Ultramaximizer, it is important that we explain some of the less obvious implications of digital audio. Once these have been explained, you will understand why Waves thought a product such as L1 was necessary and you will be a in a better position to make use of its powerful features. The operation of the L1 breaks down into two main areas:

1 - maximum level of the digital signal through proprietary peak control.
2 - maximum resolution of the signal through dithering and noise-shaping.

About maximum level

The maximum level of a digital signal is governed by the highest peak in the file. Simple normalization finds the highest peak, then raises the entire signal so that this peak is at the maximum value. However, many of these peaks may be of very short duration and can usually be reduced in level by several dBs with minimal audible side effects. Those familiar with digital editing systems may even have proved this for themselves by 'redrawing' some troublesome peaks by hand. By transparently controlling these peaks, the entire level of the file can be raised several more dB than by simple normalization resulting in a higher average signal level.

The L1-Ultramaximizer avoids the possibility of overshoot by utilizing a lookahead technique that allows the system to anticipate and reshape signal peaks in a way that produces the bare minimum of audible artifacts. Because there is no possibility of overshoot, L1 can be used with absolute confidence in situations where brickwall limiting is important.

About maximum resolution

Any digital signal processing that alters the original digital data (mixing, gain changes, EQ, dynamic processing, etc.) generally increases the number of bits required to represent the signal. Conventional truncation results in a loss of signal-resolution each time the signal is processed. The human ear uses this low level information to construct a mental image of the stereo soundstage, so any compromise in this area manifests itself as a loss of spaciousness and transparency. Waves' IDR prevents this loss of critical low-level detail.

Even when processing 16-bit signals, it is normal to process with at least 24 bits resolution, or more (as in the L1, which is now double-precision). However, as soon as the resolution is pulled back down to 16 bits by rounding or truncation (by removing the bottom 8 bits), the resulting rounding error produces an audible distortion at low signal levels, and a permanent loss of digital resolution that can never be recovered. If the audio signal is repeatedly processed and truncated back to 16 bits, the losses accumulate, causing a significant loss of fidelity, most evident as a loss of the tonal subtleties of low-level sounds within a mix. The human ear uses such low level information to construct a mental image of the stereo soundstage, so any compromise in this area manifests itself as a loss of spaciousness and transparency.

The solution is to properly dither and noise-shape a signal each time the wordlength is increased then reduced (such as nearly every digital signal process will require).

Why use them, and what are they?

Proper dithering is simply this: before the requantization (reduction of the wordlength), a precisely controlled amount of noise (termed 'dither') is added to the signal. This can convert the low-level nonlinear distortion caused by truncation into a simple steady hiss, thereby removing all traces of low-level non-linearity, but at the expense of a very slightly increased background noise. Obviously, increased noise levels are not ideal in high quality audio applications, but fortunately, the perceived level of this dither noise can be greatly reduced by **'shaping' the noise in such a way that it falls into an area of the audio spectrum where the human ear is least sensitive.**

The main point of maximum resolution is simple: to 'capture' the best possible quality into a shorter wordlength (smaller bitdepth) from a longer wordlength (higher resolution).

A bit about IDR™

IDR is a Waves proprietary noise shaping dither system developed by the late Michael Gerzon and Waves, and signifies a major advance in preserving and actually increasing the resolution of the digital signal being processed. You can utilize IDR each time during subsequent processes (for 16-bit storage), or once at the end of a high-resolution chain (say, 24-bit), in order to ensure that the final signal has the maximum resolution possible. IDR is of particular benefit when data is deliberately requantized from 24-bit to 16-bit, 16-bit to 8-bit and so on. Dithered outputs of 20, 16, 12, and 8-bit are available. The current version of L1 now features double-precision resolution, that is, that for TDM systems, all internal limiting and gains calculated with 48-bit fixed precision; on native systems, 64-bit floating precision is used. Dithering back to 24-bit output is now possible for the new DVD and other delivery media, as well as for archives of masters.

By using the IDR implemented within the L1—Ultramaximizer, optimum results may be achieved during final file preparation, mastering, and quantization or requantization. The greatest possible implementation of IDR available is in the L2-hardware limiter from Waves, and features 9th-order noise shaping (the L1 has 2nd-order).

Therefore, L1 can maximize the audio level, then increase the perceived resolution of the resulting file.

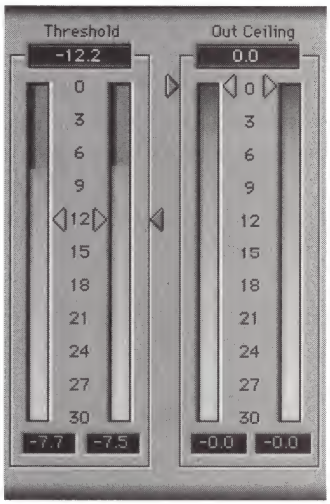

The peak limiter of the L1

The peak limiter section

1 - Open a soundfile in your host application. Select part or all of the soundfile you wish to process with the L1. Pass audio through the L1; for more information, see your host application system manual.

2 - Listen to the output of the L1. Set the threshold of the limiter by dragging the left triangle (Threshold) down. Leave the Out Ceiling setting at the maximum value: 0.0 dB.

3 - When the threshold is exceeded by the signal, you will see the Gain Reduction indicated on the single Attenuation meter to the right. Set the threshold about 4 to 6 dB lower than the peaks in the Threshold meters (input). You will see that you have from 4 to 6 dB of peak limiting indicated on the Attenuation meter (output).

4 - You'll notice that as you pull the Threshold down, the output level goes up. Leave the Out Ceiling setting at the maximum value:

0.0 dB. This is your maximum peak output.

Notice that you have significantly increased the output level. If your threshold is at -12dB, then you have pushed the signal up 12 dB (not that this is recommended)! With moderate gain reduction, the maxi mum level of a file can be significantly increased with minimal audible effect.

Only the signal above the threshold is limited; all signal below the threshold has a constant gain change that

is controlled by the difference between the Threshold and the Out Ceiling. It is this function of the L1 that allows you to maximize the level with the amount of headroom desired.

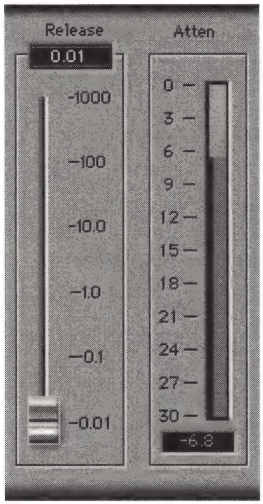

The release time and Attenuation meter

5 - Adjust the Release time to suit the application. For most sources, leave the Release time set to 1.0ms; mastering: 3 to 7ms. The release time controls how fast the L1 recovers to the constant gain after a peak is encountered.

The IDR section

The IDR processor

After reading this overview, make sure to read the chapter - **Important IDR info** - which will offer you complete IDR technical information so that you may be able to better choose from the IDR options in L1. Take the time to read about this new technology so that you may achieve the best processing for your application.

In the L1, there are two types of Increased Digital Resolution dither technology: **type1** and **type2.** The differences between them will likely influence your choice for your project.

- *type1* gives no nonlinear distortion.
- *type2* exhibits lower dither level.

It is up to you to choose no distortion or lower dither level for your soundfile. More information is in the section called **IDR Dither Type & Noise-shaping combinations.**

Basic control of IDR

- Select Quantize level for desired output (20, 16, 12, 8 bit, or *no IDR*) by clicking the button to toggle through the choices. Alternately, if you option-click-hold on the Quantize button, a popup menu will allow direct selection of bitdepth.
- Select Dither (type1, type2, non) by clicking the button.
- Select the type of noise-shaping (Moderate, Normal, Ultra, non) by clicking the Shaping button. (The following chapter has extensive information on the combination of these 3 options.)

For optimum results, level maximization (and at least normalization) should also be done! This is why the L1 includes both advanced peak controlling and IDR together: one step maximizes both.

An example of IDR resolution improvement

To hear what truncation (an extremely severe form of quantization error) sounds like at the 1-bit level (present in every signal, no matter the bitdepth) and how Waves' IDR and Noise-shaping works, try this simple (and exaggerated) example:

NOTE: L1 is currently optimized for 48, 44.1, 32, and 22kHz files.

1 - Using a 16 bit, 44.1kHz soundfile, select a desired region. Select the +L1 plug-in within your application and set your system so you can monitor the L1 output.

2 - Set Dither Type to non, and Shaping to non. Set Quantize to 8 bit for easier auditioning of quantization distortion.

3 - Listen to the output of the L1 (preview or realtime, depending on your host application).

4 - To more easily hear the quantization error, reduce the the input level at least 30dB, until you hear the sound becoming distorted in a 'fuzzy' way. If necessary (depending on the input level), pull the Out Ceiling level down also, to about -15dB. Leave the Threshold at 0.0. (You will have to raise your monitor levels quite a bit to hear the distortion noise - so do NOT click Bypass until you turn your monitor level down!) This quantization distortion you are hearing is present in varying degrees in ALL digital signals at the 1-bit level.

5 - Click Shaping twice. It changes to Normal. You will hear the music cleanly (the nonlinear distortions are gone), and that a steady hiss in the background is now present. Of course, since you have reduced the input level, this hiss will be much more prominent than in actual work and is exaggerated just like the distortion. You can cycle throughout the Shaping options to listen to the quantization error-removing effects of IDR's different noise-shaping filters.

6 - For this example, leave Dither set to non. Audition the noise-shaping settings to familiarize yourself with the distortion-removing properties of IDR. For 16-bit work, the flexibility of the IDR implementation of L1 gives you the choice of dither and noise-shaping. Please read the following sections on type1, type2, and Noise-shaping. For specific applications, see the appropriate chapter (8-bit, ≥16-bit, etc).

Type1

This is the 'purist' technology. It is designed for no nonlinear distortion or modulation noise at low levels and combines optimal dither noise with psychoacoustic noise shaping.
When signals are subjected to several stages of higher-resolution processing and re-quantizing back to lower bitdepths, the design of resolution-enhancement must satisfy different requirements than a technology designed just for one-stage use. If applied several times in succession, a digital resolution enhancement technology optimized for one-stage CD mastering can produce unwanted side-effects.
Waves **type1** technology, however, is the first optimized for use at every processing stage, allowing for the effects of cascading and subsequent signal processing. Type1 is also optimized to cause minimal side effects when used with stereo signals.

Type1 is the recommended choice for use with 20 and 16 bit files processing and other high quality mastering applications. By combining level maximization (peak controlling) and IDR processing, 16 bit files created from 20 or 24 bit masters this way can have an apparent resolution of 19 bits — an 18dB improvement.

Type2

Type2 also uses dither with a similar noise-shaping curve, but the dither is of a unique kind designed to minimize the amount of noise added, thereby giving a lower noise level than the IDR type1 process, but at the expense of some low level distortion.

Type2 does have some advantages for high quality mastering as well, and it is purely your choice whether the ultimate in low distortion of type1 is preferred, or the additional reduction in dither level of type2. It may be used for multimedia applications of all 8-bit/44.1kHz files and many 8-bit/22kHz files, dependent on source material. Type2 is "black" with no input signal, that is, there is no dither signal if there is no input. See the 8-bit multimedia mastering chapter for more information.

Noise-shaping options

Another way to decrease the perceived amount of noise and increase perceived resolution is to 'shape' the frequency content of the noise so it matches the ear's sensitivity curves. In basic terms, noise-shaping shifts the energy of the

noise to the frequency ranges where we hear it the least.

The three options of Noise-shaping provided in L1's IDR section push more of the noise energy to higher frequencies above 15 kHz, where our ears are least sensitive, and reducing the noise energy of lower frequencies. The three Noise-shaping options differ in the amount of this 'shifting action'.

Moderate is provided for 8/12 bits files, but also is completely usable for 16 bit files.

Normal is the recommended option for use under most conditions for all bitdepths.

Ultra is a very high-quality setting, suitable only for use at the very last stage of mastering high-resolution files (16-bit and higher) targeted for high-quality digital media. It is theoretically possible that the relatively high amount of high frequency energy could cause undesirable side effects *if the signal is going to be processed or digitally edited again.* Therefore, it is best that Ultra is used in the last stage of file preparation. However, with many thousands of L1-processed products, no such situations have been reported. Since it is *theoretically* possible, we wish to inform you of it. Most companies do not. These theoretical side effects might cause clicks at later edit points if 'Ultra' noise shaping was used, *if played back on poorly designed D/A converters.* However, state-of-the-art designs rarely have these troubles anymore.

Of course, the effect of Noise-shaping is even greater when used with type1 or type2 dithering, since Noise-shaping reduces the audibility of the added dither noise.

Now try the full effect of IDR technology by listening to the same material, with both IDR types and different kinds of noise shaping. The most obvious places to examine are notes or reverb during the end of the sound, or 'tail'; it is during this time that quantization error is most audible, although it is present on all low-level signals (such as elements that are soft in a mix, etc). Since the entire issue of dithering is a very subtle one, we recommend you listen to a rather long piece of audio (2-3 minutes) of high-quality, say 20-bits if possible, with a good dynamic range. Jazz and classical recordings are ideal.

If you don't feel you fully understand the trade-offs between IDR and Noise-shaping settings, the option that will generally work well for CD-mastering is type1 with Normal noise-shaping. For minimum noise with 16-bit and greater files, type2 Ultra; maximum resolution, type1 Ultra.

IDR™ technology was designed by **Michael Gerzon**, a Gold-medal AES fellow, and a world authority in psychoacoustics. He also invented the SoundField microphone, and was the major contributor to Ambisonics™. The design of IDR is a result of his long-term researches, dating back to 1982, with many of the other leading experts in digital resolution enhancement technologies.

Chapter 4 - Important IDR Info

In L1, Waves provides users with three dither options:

1 - No dither <non>. This is normal truncation and gives a high degree of nonlinear distortion at low levels. This is what most signal processors do.

2 - IDR type1 dither. This dither adds a certain amount of noise, causing a 5 dB increase in background noise compared to no dither, but completely eliminates all low-level distortion and signal-dependent modulation effects. The result is a very transparent and clean low-level sound, with a very high resolution.

3 - IDR type2 dither. This dither adds virtually no audible noise, and so is 5 dB quieter than **type1**, but it still gives some low-level distortion. However, this distortion is generally much lower than with no dither at all.

There are four Noise-shaping options:

1 - **None**. No noise shaping, resulting in high audible hiss levels, and high distortion levels when used without dither.

2 - **Moderate**. This typically reduces perceived noise by around 6 dB, and slightly reduces audible distortion when no dither is used.

3 - **Normal**. This typically reduces perceived noise by around 8.5 dB, and somewhat reduces audible distortion when no dither is used.

4 - **Ultra**. This gives the greatest perceived noise reduction, typically 10.5 dB.

The noise shaping options of Waves IDR differ from those in most commercial noise shapers by avoiding shapes of noise that sound "colored" to the ear. While this gives less perceived noise reduction than in some of the most "extreme" shapers on the market, the resulting sound is more pleasant, less colored, and takes more kindly to several generations of processing of sound files.

The noise reduction figures given here apply to sampling rates of 44.1 or 48 kHz. At the lower sampling rates of 22 or 32 kHz, the noise shaping options are re-optimized for best results at the lower sampling rates and give smaller, but still useful, reductions of perceived noise.

If audibility of noise were the only factor, the choice would almost always be to use Ultra noise shaping, but in some situations, heavy noise shaping of the Ultra kind can have some disadvantages, and the alternate settings, such as Normal or Moderate may be better.

As mentioned previously, it is theoretically possible for any extreme noise-shaper to cause trouble with cheap DA converters. Therefore, we want to include these small notes.

Ultra shaping should be avoided in the following situations:

(a) Subsequent digital editing, when the signal is subjected to later editing. At the edit points, an extreme noise shaping might cause audible "clicks", which may be disturbing. An example application in which you would avoid use of Ultra shaping would be on CD's with production music or sound effects that would certainly be subject to further digital editing.Use Normal for optimum compatibility.

(b) Poor Error Correction, i.e. when a signal is destined for a carrier medium with poor error correction, such as CDs pressed in pressing plants with poor quality control (which are not so widespread these days). When errors that are not properly corrected occur, the Ultra setting, like all forms of heavy noise-shaping and other resolution enhancement technologies, might cause audible background crackles, but most especially on very cheap CD players. While these effects generally don't occur on the majority of mid- or hi-fi CD players, they can be noticeable on very cheap products. The amount of such crackles on poor pressings is greatly reduced by the Normal noise shaping.

(c) If high-gain treble boost is subsequently employed. This can cause the strongly boosted higher frequencies used by Ultra noise shaping to become so high in level that they become unpleasantly audible, or to feed excessive noise energy into loudspeakers. Therefore Ultra shaping is best avoided in situations where subsequent heavy equalization may be used. A much lower boost of high-frequencies is used with the Normal and Moderate noise shapers. Ultra will not affect Dolby or broadcast encoders.

Recommended IDR settings

All of the Factory Presets (in the Load menu) are relatively self-explanatory. Keep in mind that any combination of dither and noise shaper can be used, but the following settings are particularly recommended for different applications and [bitdepths]:

General Purpose high-quality use, including material liable to be edited and EQ'ed: **type1-Normal** [24, 20, 16, 12]

Lowest Noise (CD): type2-Ultra [24, 20, 16, 12]

Low Noise/Highest quality: type1 - Ultra [24, 20, 16, 12]

Low noise while allowing editing/EQ: **type2 - Normal** [24, 20, 16, 12, 8 - excellent for 8/44.1 files]

High Quality, with lowest risks of spurious noises on edits or cheap CD players: **type1 - Moderate** [24, 20, 16, 12, 8 - excellent for 8/44.1 files]

Low noise, with lowest risk of spurious noises on edits or cheap CD players: **type2 - Moderate** [24, 20, 16, 12, 8 - excellent for 8/22kHz files]

Lowest Noise for Multimedia (ideal for voice): **non - non** [8, excellent for 8/22 or 8/11 files]

Lowest Distortion for Multimedia music (good for music or continuous sound): **non - Normal** [excellent for 8/22 files; not for use with 8/11]

Chapter 5 - 16-bit (and higher) mastering

Here are the basic steps of using L1 in a 16-bit, 44.1/48kHz application. These steps also apply to 24- and 20- bit mastering.

- All processing, EQ, sample rate conversion, dynamic changes, etc. MUST be done before L1 processing. The L1-Ultramaximizer should be the last processing of the file. Ideally, dithering occurs only once.

- If the left and right channels need to be balanced, adjust the Input levels (separately or together) by using the input faders.

- Using a 16 or higher bit input file, set the Threshold for desired peak limiting. For suggestions on how much limiting to do for certain applications, see the Specific Applications chapter. In general, set Threshold for about 4-6 dB of Gain Reduction in the Atten meter.

- Now take the Output Ceiling up to the maximum peak output you desire. You can take this Output all the way to 0.0 dB without any clipping. For CD's, a recommended setting is -0.3dB; for more information read the Peak Clipping section. Factory Presets already set the Out Ceil control to the recommended value.

- Leave the Release time set to the default (**1.0** millisecond).

- Set Quantize output for **16-bit** (for CD/DAT; or 20, 24, for higher archival or mastering medium if your hardware supports the transfer of 16+ bits).

- Set Dither type (type1 or type2). **IDR type1** is recommended for most high-resolution applications.

- Set Shaping (Moderate, Normal, Ultra, non). **Ultra** and **Normal** are recommended for most high-resolution applications.

- Select Digital or Analog Domain mode. Analog domain is recommended for all final production masters, as it gives a little more protection from clipping in poorly-designed DA converters without compromising any L1 processes. For more information see Digital/Analog Domain paragraph in the IDR controls chapter.

Chapter 6 - 8-bit multimedia mastering

There is no single best way to process 8-bit files because each method has its strengths and weaknesses. The best approach for any given project depends both on your tastes and on the use to which the finished audio file will be put. In order to provide some practical guidelines, three methods are described here which have been formulated by top beta-testers and multimedia developers. These techniques employ peak-limiting, with optional combinations of dithering and noise-shaping.

It is essential that your host application be able to write a physical 8-bit file on disk (we assume you're working from 16-bit files; there is very little improvement trying to use L1 on existing 8-bit files). Consult your host application manual for details, or their tech support staff.

First: do all other processing, including Sample-Rate conversion

Second: choose IDR settings

For lowest noise

Regardless of sampling rate, for the lowest level of background noise when working with an 8-bit file, it is recommended that you use the L1's peak limiter but without dithering or noise shaping. The appropriate IDR settings are: Quantize-8-bit, Dither-non, Shaping-non.

For higher resolution

When working with 22kHz sampling rate music files, it is generally best to use Ultra noise-shaping and peak limiting without dither (type1 or type2). The corresponding IDR settings are: 8-bit, non, Ultra.
This processing approach yields good results for continuous signals, such as music-only or background files.

For highest resolution

When working with 8-bit files with sampling rates of 22kHz or greater, you can use Ultra noise-shaping with IDR type2 dithering. The appropriate IDR settings are: 8-bit, type2, Ultra. However, use will depend on the source material.

Third: reduce the bitdepth

Some applications require specific settings to write smaller bitdepth files. Read your host application manual or contact their tech support staff for information. The following observations are relevant when the bitdepth of a file is being reduced–for example, if you have a 16-bit file and need to create an 8-bit final output, or if you have a 24-bit file and need a 16-bit final output.

Chapter 7 - Use the L1 last

- It is recommended that L1 be used as the final process after all dynamic and EQ adjustments have been made.

- Only when all these processes are finalized should the question of peak level be addressed. Instinctively, it might seem appropriate to Normalize the file once all other processing has taken place, but in practice, it may be better to set the peak levels to around 1dB below clipping using L1. For an explanation of this reasoning, see the notes on digital clipping.

- The choice of IDR setting depends on the final use to which the file will be put. Type1 or 2, Normal, is recommended for most work. Type1 or 2, Ultra, is considered best for final mass production of 16bit and greater masters, and for producing a complete disc master which will undergo no further edits. An example of this would be a production master DAT run off from a hard disk editing system in a single pass and where no further editing is anticipated (all timing is finalized). This DAT would then be transferred unchanged through to the glass master.

- When using programs such as Jam, MasterList, or MasterList CD which in effect string together separate regions of the soundfile, you must perform destructive processing with the L1. Since MasterList/CD is a form of editing (playing regions/songs with silence or splices between audio regions/ songs), it is recommended to use Normal noise-shaping with either Dither type1 or type2. However, thousands of L1-processed masters have been produced with type1 Ultra with no unusual problems.

- If you must process or EQ any previously L1-processed file, you will need to create headroom by lowering input on those later processes, then probably re-limit to restore the average level.

Chapter 8 - Important notes on digital clipping

The digital words representing an audio signal at each moment have a maximum possible positive value and a minimum possible negative value defined by the bit depth of the file format. Any attempt to force an audio signal beyond these maximum permitted values, for example, by applying excessive gain, will result in the audio signal being clipped. Clipping distortion generally sounds quite unpleasant and is to be avoided.

However, there are other ways in which a signal can become clipped, and some of these are far from obvious.

Peak-normalized signals

A '**Normalize**' process allows a file to be processed in such a way that the maximum peak level within the file just reaches (but does not exceed) the digital zero or clipping point. This is obviously desirable as it means that the file is as 'loud' as is possible without clipping, and in turn, this maintains the best signal-to-noise ratio, especially at low bit word lengths.

In situations where a higher average sound level is required, L1's peak limiter allows the typical level of signals to be even further increased by gently pulling down the gain, without audible nonlinear distortion, of waveform peaks. L1 can simultaneously rescale the audio data so that the limited peak signals approach or just reach digital zero.

However, by storing soundfiles at the maximum possible level, there is a risk that any subsequent processing may take these peak levels too high, resulting in clipping distortion. Peak limiting to 0dB, by whatever means, leaves no margin for any subsequent increase in peak level.

Intuitively, you might think that simple gain reduction could be applied without incurring the risk of clipping, and equally that any increase in gain would be sure to cause clipping. This is true. You might also think that applying an EQ boost at any frequency might result in clipping depending on the peak energy level within the band being equalized. Again this is true.

What is far less obvious is that applying an **EQ cut** <u>also</u> runs the risk of causing clipping. To prove this would take a lot of math, but the following description should help get the point across.

At any instant, the peak level of a signal may be the result of several components at different frequencies and at different phases relative to each other. Some components will add while others will subtract, but what happens if you 'EQ out' a frequency that would otherwise be subtracting from the peak level by virtue of its phase? The peak is now *higher* than it was.

For most audio material, this effect will be relatively small, typically increasing peak levels of the order of 0.3 dB, but is possible that under unfavorable circumstances or with nontypical signals, the peak level could increase significantly more than this.

If an effective peak limiter like the L1 is used, because it forces the signal to skim the peak level more often, this likelihood of clipping in this way is further increased.

In practice, filters attenuating mid or high frequency components tend to cause the kind of increase in peak level

described, but high pass filters that attenuate the bass can sometimes cause much larger increases of peaks of the order of several dBs on heavily limited signals. The phase response of certain high or low-pass filter types can also increase peak levels by up to 4 dB or thereabouts.

Bearing in mind what's just been said, it might seem logical to keep the signal peaking a few dB below digital zero until all processing has been carried out. After that, you can safely normalize the signal–or can you?

A related problem with peak clipping can arise when a normalized soundfile or signal is converted to a new sampling rate. The reason has to do with the sample-rate conversion process itself, and during sample rate reduction, the signal is effectively being filtered; the available audio frequency range is smaller at lower sampling rates. Such filtering can increase peak sound levels in exactly the same way as attenuating equalizers can.

But even when increasing sampling rate, an increase of peak level can occur. This is because the continuous-time audio waveform is represented in the digital domain only by its values at the sampling instants. It is perfectly possible for the peak value of the continuous-time audio waveform to occur at instants lying between two sampling instants, and thus to be higher than the peak value at any of the sampling instants.

When changing the sampling rate, new sampling instants are chosen for the continuous-time audio waveform, and these new sampling instants may coincide with an increased peak lying between the original sampling instants. This is especially likely to occur with signals with a lot of high frequencies, since such signal waveforms change more rapidly between the sampling instants.

Though artificially contrived signals can be created to really show up this problem, in real life an attenuation of at least 0.3 dB or so prior to conversion should provide adequate protection against clipping. You might expect sample-rate-converter designers to account for the possibility by designing in a small amount of attenuation, and the cheap ones generally do not.

But can you safely normalize a file that you know is at the final sample rate? Unfortunately not, because many compact disc players (and some other digital consumer equipment) use oversampling digital-to-analog converters (DACs) to produce the analog signal fed to the amplifier. Such oversampling converters involve a sampling rate conversion process which can (and does!) cause audible peak clipping. Once again, some designers appear to have overlooked this problem, although not as widely as they did in earlier DAC designs.

Practical Precautions

L1 incorporates an 'analog domain' processing mode that estimates the between-sample peak levels to reduce the clipping risk, but even here, the process is not foolproof, so that some headroom to avoid clipping may still be needed.

If you want to normalize and avoid clipping, normalize your files after all other processing, including sample rate conversion. However, if you're going to use L1 limiting, normalization may be completely redundant. In all cases, the L1 should be used last.

Chapter 9 - L1 Controls descriptions

Quantize

This controls the final bit depth (8, 12, 16, 20, 24) of the output of L1. It is not related to the input bitdepth.

16-bit

If you are saving the final output to DAT or CD recorder, then set the Quantize to 16 bit; Waves plug-ins work at an internal resolution of 24 bits or higher during all processing. The Quantize setting "captures" the best data possible for the selected bitdepth.

8- and 12-bit

If you are saving the final output for low-resolution multimedia soundfile work, then you should set the output to the appropriate bit depth (8 or 12 bit). The file should be sample-rate converted PRIOR to level maximization and dithering.

20-bit

If you are saving the final output to 20-bit storage medium, set the Quantize to 20-bit.

24-bit

New feature! Now the L1 works at 48-bit resolution, and is dithered back to 24-bit when this option is chosen. Ideal if you have a 24-bit archival medium or intend on a 24-bit file bounce for later work. In previous versions, the signal was not dithered with this selection, and in very old versions of L1, the selection actually said "no IDR".

Dither

This controls the type of dithering process (**type1**, **type2**, **non**). Please read the Important IDR Info section completely for better comprehension of the options that the two modes of IDR present.

In brief:

IDR dither **type1** yields lower distortion (greater linearity);
IDR dither **type2** yields lower dither level.

Shaping

Noise-shaping may be selected to suit your preference for each application. Select from (**Moderate, Normal, Ultra, non**). Normal and Ultra are recommended for optimal result with type1; Moderate and Normal are recommended for optimal result with type2, however there are no specific restrictions! Your preference is the most important factor.

Noise-shaping may be used without IDR being engaged, although it is not an optimal use of the technology. For multimedia work, it is a very good option for music files. For complete information, read the Important IDR Info section.

Digital/Analog Domain

When in the Digital Domain position, absolutely no sample will be over the Out Ceiling value. However, after analog conversion, it is possible to have peaks higher than in the digital domain. This is due to very complex digital audio issues involving peaks "between the samples". Almost all quality-made digital-to-analog converters have at least 3 dB headroom to allow for these peaks; many have at least 12dB headroom.

You would want to use the Analog Domain position when you wish to have absolute control over any peak that occurs in both the analog and digital domains. Some examples (as mentioned previously in the manual): you wish to accommodate poorly designed DACs; the file will undergo further modification, such as ADPCM data-reduction; or, you wish to have a signal that can be broadcast without further peak controlling. In these cases, brickwall limiting is desired in both domains, and you should put L1 in the Analog Domain mode.

Chapter 10 - Recommended settings

These are intended as a few guidelines for use of L1. Of course, your preferences will be based on your experiences; the settings here are suggested just for quick reference or as beginning points for your session. The **Important IDR Info** section contains critical information for you concerning all of the choices in the use of IDR technology. Sample rates would include pull-down, pull-up, and other rates close to those mentioned here.

High-resolution mastering (no further digital editing)
Setting: Dither type1, Ultra, 4 to 6 dB gain reduction, default release time.
Bitdepth: 16, 20, 24.
Sample rates: 48, 44.1, 32kHz.

Multimedia and low-resolution mastering
Setting: Dither type non, Moderate/Normal shaping for music, shaping non for voice-only; up to 15 dB gain reduction, (this heavy gain reduction may need longer release times, such as 20-40ms.) Low peak/rms ratio sounds (such as synthesizers) should not have very heavy gain reduction with short release times.
Bitdepth: 8 and 12.
Sample rates: 22 and 32kHz, and similar rates.

For 11kHz and below
Set Dither to non; set Shaping to non. Use only the peak limiter with 8-bit Quantize setting.

C1
Parametric Compander

Table of Contents

Chapter 1 ..... About the C1 5

The C1 in brief 5
C1 Component plug-ins 6
C1's Setup Library 6

Chapter2 ..... Comp/Exp module 7

About the compressor/expander 7
Controls 8
Bypass control 8
Level reference control 8
Makeup gain 8
Threshold 8
Ratio 9
Attack 9
Release 9
PDR (Program Dependent Release) 9
EQ Mode switch 9
Gain reduction Meter 10
Comp/Exp control level meter 10

Chapter 3 ..... Gate/Expander module 11

About the Gate/expander 11
Gate/expander module controls 12
Bypass control 12
Gate/expander control 12
Floor 12
Gate Thresholds 12
GateOpen 12
GateClose 12
Attack 13
Release 13
Hold 13
EQ Mode switch 13
Gain reduction Meter 14
Gate/expander control level meter 14
Linking controls in the two dynamics modules 14

Chapter 4 ..... The filter module .. 15
About the filter module 15
Filter Controls 16
Type 16
Frequency 16
Q 16

Chapter 5 ..... Other controls and I/O metering .. 17
Lookahead switch 17
About lookahead 17
Input/Output Graph 18
About the input/output graph 18
Indicators and controls 19
Comp/Exp display 19
Gate/Exp display 19
Keying mode 19
Key modes 19
Monitor switch 20
Audio 21
Sidechain 21
Passive (unprocessed) 21
Input and Output Levels 21

Chapter 6 ..... Basics of operation .. 22
The Compressor module 22
LowRef/PeakRef mode 24
Which level reference mode should you use? 24
Adjustment of the compressor in LowRef mode 25
Adjustment of the compressor in PeakRef mode 25
How to adjust PDR 45

Chapter 7 ..... The Gate/Expander module .. 26
To use as a Wideband Gate 26
To use as a Wideband Expander 27
To use the Floor control 27
The Filter module 28

Chapter 8 ..... Sidechain Tutorial .. 29
Next step 29
Another de-esser 29

Chapter 9 ..... Intro to Combination setups .. 31
Classic Combinations 31
High-level Compressor + De-Hisser 31
Compress + De-reverb 33

Chapter 10 ..... Setup Library listing .. 35
Simple Setups 35
Compression 35
Noise Reduction 35
De-Essers 35
Enhancers 35
Keying Setups 36

Chapter 11 ..... Conceptual block diagrams of all EQ mode options 37
Key modes 40

Chapter 12 ..... C1 Technical Specifications .. 43

The control window of the Waves C1 Compressor/Gate

Chapter 1 - About the C1

The Waves C1-Compressor/Gate is a studio-quality stereo dynamics processor and a "parametric compander". The C1 is also a fully featured dynamic filtering processor capable of many unique effects including sound sweetening, enhancement applications, and applications requiring correction of sound quality problems. For maximum flexibility and efficiency, the C1 breaks into smaller components, such as the C1comp and C1gate, or the C1c/sc (comp with sidechain). This lets you choose exactly the dynamic tool needed for the job you have. The C1comp and C1 gate (which don't have sidechain or lookahead) now can be used at 88.2/96kHz on some native platforms.

The C1 is designed for optimal subjective sound. While it is rich in features and functionality, and can do many things unavailable on any other dynamics processor (whether analog or digital) on the market, the design is based on features that sound good to the ear.

The C1 - in brief

The C1 consists of three independent stereo processing modules that can be used together in many different ways:

- Compressor/expander module
- Gate/expander module
- Filter/equalization module

The compressor/expander module (**Comp/Exp**) is capable of operating as a high-level compressor, a gentle high level expander, or as a mid-level compressor whose compression can be "tuned" at any user-chosen signal level. This module also has a unique cancellation dynamics mode for cancelling out very high-level sounds.

The gate/expander module (Gate/Exp) is capable of operating as a fully-featured expander or gate. It can also be used to take lower level signals and "compress them upwards".
The filter/equalization module (Filter) has a variety of filtering capabilities for many different kinds of processing, from conventional sidechain dynamics processing to a unique bandsplit dynamic EQ mode. This bandsplit configuration allows compression, expansion, or gating of any desired frequency band.

Each module can be set for different operating modes, plus various combinations can be used for efficient use of your processing power. For example, the Comp/Exp can perform wideband dynamic processing while the Gate/Exp module gates only the high frequencies (hiss gate). Another example would be using the compressor for bass compression while the Gate/Exp module gates rumble.

C1 Component plug-ins

On most host applications and platforms, you can select just the audio processing you need and use only the power necessary to do the job. The C1 has several plug-in components in the menu, and most have mono and stereo options.

Each has their own separate Factory Presets and User Presets. Information on each of the Factory presets will be in the C1 documents installed onto your hard disk. For more information about Component plug-ins, read the Component chapter in the WaveSystem Manual.

As of this printing, here are the current components; new ones may be added in future updates, and will be noted in the plug-in Read Me file. Some examples for each component are listed.

- **C1comp** — a wideband compressor/expander, no sidechain; *(voice or instrument, individual track compression, gentle soft-knee mastering, 'uncompression' by soft-knee upward expansion, soft-knee limiting)*
- **C1gate** — a wideband gate/downward expander, no sidechain; *(track gating, smooth voiceover downward expander, gate hold, low-level compression)*
- **C1c/sc** — compressor with sidechain or split-band compression; *(HF limiter, LF limiter, de-esser, de-popper, sidechain drum gate, any-Frequency limiter/compressor/expander, no-pump sidechain mastering compressor)*
- **C1comp/gate** — compressor and gate, wideband, when sidechain is not needed; *(voice or instrument, general "outboard" application, track insert)*
- **C1c/g** — the "full" plug-in.

The Setup Library

For extensive control and fast setups, the C1 system includes an extensive Setup Library. It will take you directly to the essence of the C1. You can also use it as quick starting points in your work and as a tutorial for the features of C1. Waves strongly encourages you to try the Setup Library. A list of all setups in the library is included in Appendix A at the end of this User's Guide.

Included in this book is the C1 Setup Guide which consists of extensive documentation written by Michael Gerzon, including his highly insightful tips on every aspect of the C1. If the C1 seems a bit daunting, check the Setup Guide for a more application-specific approach.

Chapter 2 - About the compressor/expander

In the menu, you'll have a "C1comp" option, which is the Comp/Exp module by itself, with no sidechain control. The EQ Mode button is omitted when C1comp is used.

The Comp/Exp module is a soft knee variable-ratio processor which at maximum compression becomes a soft-knee limiter. The softness of the knee has been designed to reduce limiter gain about 3 dB at the limiting threshold.

The dynamic processor is unusual in two respects. First, for moderate compression ratios, the actual dynamic curve compresses over a range to about 20 or 30 dB above threshold, rather than extending the range indefinitely high. As explained below, this is much more flexible than a conventional compressor, while still allowing conventional use when working within 20 or 30 dB above threshold.

Second, an extended ratio control range permits (1) controlled high-level expansion by using ratios of between 0.5:1 and 1:1 and (2) a unique high-level "cancellation" mode beyond limiting, displayed as a *negative* compression ratio. Very high level signals well above threshold can be nearly cancelled out completely by using these *negative ratios.*

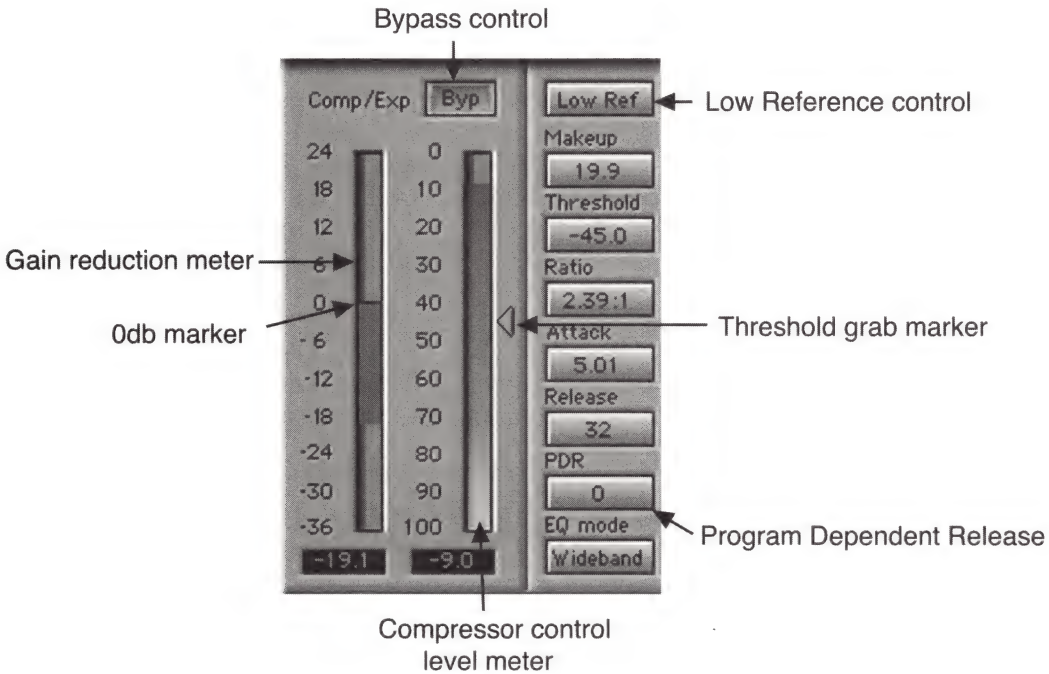

Controls

Bypass control

This switches the module's processing in and out. Shows **Bypass** when the module is out.

Level reference control

This has no effect on the sound but affects the output level and how the makeup control works:

LowRef mode [Low-level reference]. This is the conventional mode for compressors, in which the output drops as the threshold is lowered (and more gain reduction occurs).

PeakRef mode [Peak-level reference]. This permits easy adjustment of the Comp/Exp when you want to keep peak levels approximately the same. For example, when in PeakRef mode, as you lower the Comp/Exp threshold, the output gain increases, keeping the output level approximately the same.

Makeup gain

This is the traditional "output" level of the Comp/Exp module and is calibrated in dB, ranging from -40 dB to +40 dB.

Threshold

Calibrated in dB below 0dBFS (Full Scale digital). This is the input level above which the soft knee compression/expansion starts acting to a significant degree. Adjust the **Threshold** by clicking and dragging on this Value Window, by dragging the triangle-shaped **grab marker** next to the Comp/Exp control level meter, or by dragging the upward-pointing **grab marker** under the input/output graph. (See illustration on next page.) You can change the Threshold by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

click and hold on any Value Window, then drag up/down to change value

This grab marker is another way to control Threshold. Drag left or right.

Ratio

This control adjusts the compression (or expansion) ratio in the region above threshold. A ratio of 1:1 corresponds to no compression, whereas approaching 50:1 corresponds to a soft-knee limiter. Ratios between 0.5:1 and 1:1 give high-level expansion. Ratios in the control range above 50:1 are shown as negative values between -50:1 and -5:1. This mode is termed "cancellation" mode and is useful for cancelling out unwanted very loud signals while retaining quieter ones. Cancellation mode can result in dynamics that sound very odd, but it does have a number of uses explained later in the manual. You can also change the Ratio by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

Attack

Calibrated in msec. Adjusts the attack time of the Comp/Exp with time constants from 0.01 msec to 1 second. You can change the Attack by clicking and dragging on the Value Window, by direct numeric entry, or by the up/down arrow keys.

Release

Calibrated in msec. Adjusts the release time of the Comp/Exp with time constants from 1 msec to 10 seconds. You can change the Release by clicking and dragging on the Value Window, by direct numeric entry, or by the up/down arrow keys.

PDR (Program Dependent Release)

This control allows the release time to vary according to program content. When set to the lowest value, release time is fixed at the value shown in the release window. You can change the PDR by clicking and dragging on the Value Window, by direct numeric entry, or by the up/down arrow keys.

At higher settings, release time is shortened for short-duration transients; a slower release is used for longer duration signals.

The PDR calibration in milliseconds is a rough indication of the length of transients below which release time will be shortened. Thus for a PDR of 150ms, transients shorter than about 150ms will have a shortened release time. For transients longer than 150ms, the release time is the one specified by the **Release** control.

EQ Mode switch

The EQ Mode switch is only available when sidechain is used and therefore is not seen in the "C1comp" component interface.

This has three modes: *Wideband, Sidechain* and *Split.*

In **Wideband** *mode*, the Comp/Exp operates as a conventional wideband compressor.

In **Sidechain** *mode*, the sidechain audio signal used to control the dynamics is first passed through the filters in the EQ module to make the dynamics respond only to a defined frequency band in the audio.
In **Split** *mode*, the sidechain audio signal is filtered, but the main signal is also split into two complementary bands: the

active band (whose frequency response is shown in red on the graph) and the *passive band* (shown in blue on the graph). Only the active band is subjected to the Comp/Exp, and the passive band is not dynamically processed in any way. By this means, a dynamic equalizer, whose action is confined to the active band, is obtained.

Gain reduction Meter

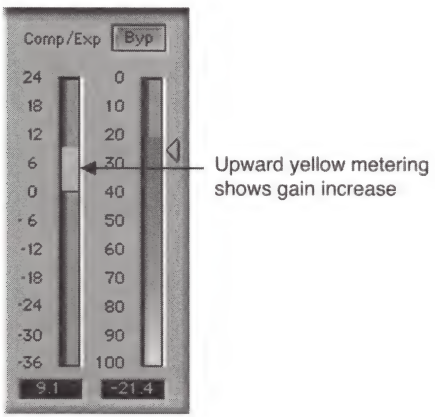

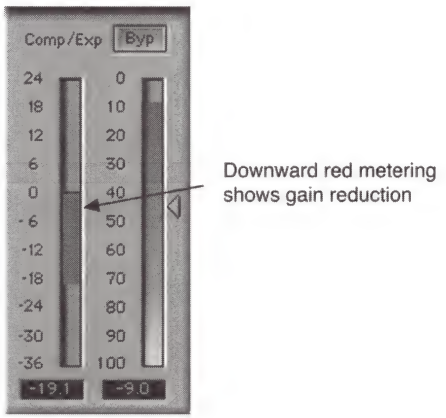

This shows the instantaneous gain reduction in dB (in red below 0 dB) or gain increase in dB (in yellow above 0 dB) of the Comp/Exp. The meter does not show the effect of makeup gain.

The value window above each meter in C1 shows the largest absolute peak since the last 'reset' of the meter. To reset any peak numeric meter, click directly on the peak value itself, or its meter bar.

> **Note:** When in PeakRef mode, the Gain reduction meter will show the *actual gain change for the signal.* This may be a bit confusing at first. For instance, if you are compressing a signal while in PeakRef mode, the peaks will be kept about the same if Makeup is set to 0.0dB. The Gain meter will show a *positive value,* indicating the amount of the gain increase necessary to keep the peaks about the same. Conversely, if you are expanding, then the meter will show how much the gain is reduced to keep the peaks about the same.

Comp/Exp control level meter

This blue meter shows the *sidechain* signal level that controls the Comp/Exp. The meter time constants match the Comp/Exp attack and release times. The threshold can be conveniently adjusted by dragging the small triangle **grab marker** beside this meter.

Chapter 3 - Gate/Expander module

About the Gate/expander

In the menu, you'll have a "C1gate" option, which is the Comp/Exp module by itself, with no sidechain control. The EQ Mode button is omitted when C1gate is used.

The Gate/Exp processor is operable in two modes: (1) as a soft-knee low-level expander for gentle reduction of low-level signals, or (2) as a gate including hold and hysteresis adjustments for well-behaved gating of low-level signals.

The Gate/Exp has a full range of adjustments which are independent of those in the Comp/Exp module, although, as explained later, some of the controls of the two modules may optionally be linked together for ease and speed of adjustment.

The extended floor control range is unique to this gate/expander. It permits not just a controlled degree of gain reduction below threshold, but even an increase in gain for low-level signals, which can create a low level compression effect when needed. This is the low-level compressor mode. In addition, Negative Floor values (indicated by a capital 'N' after the value) are for gate cancellation mode. This is useful in *split modes* for single-ended dynamic noise reduction, explained later in this manual, and even more extensively in the C1 Setup Guide.

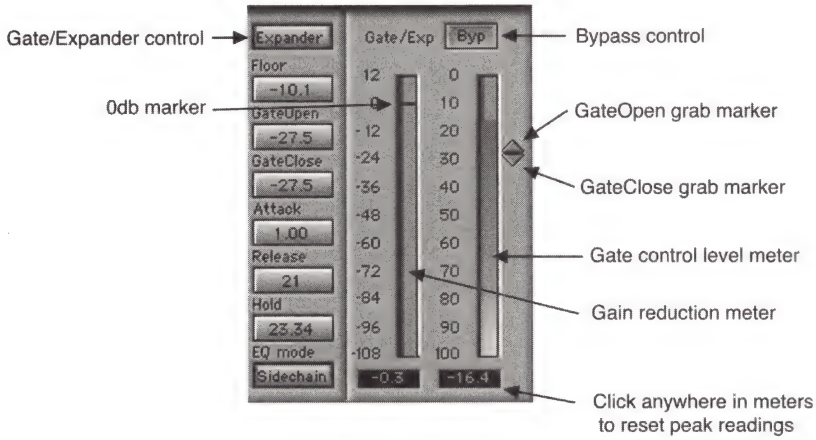

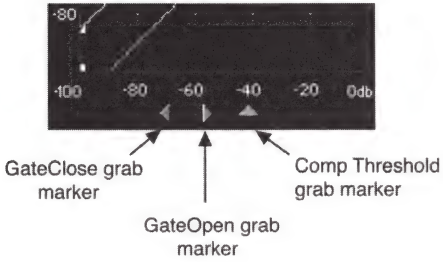

Gate/expander module controls

Bypass control

This switches the module's processing in and out. Shows **Bypass** when the module is out.

Gate/expander control

You can switch the mode of the Gate/Exp module to either *Gate* or *Expander.*

Floor

Calibrated in dB. **Floor** adjustment allows some signal to remain in order to avoid a complete signal cutoff. This adjusts the amount of residual signal gain at low levels when the gate is closed or the expander is at large gain reductions. When **Floor** is set to *-infinity dB*, essentially no signal is left at low levels. When set to 0 dB, the signal is unchanged at low levels. When set to positive gains, up to 12 dB, the signal level at low levels is actually increased, creating a low-level compression effect whereby quiet signals are increased in level. This low-level compression effect is particularly useful in the soft-knee expander mode.

In expander mode, the **Floor** control also goes beyond the *-infinity* setting to negative polarity settings (-100 N to -10 N). This is indicated by a negative number of dB followed by the letter N to indicate "negative polarity". This gives the expander a steeper expansion ratio below threshold, and of cancels sounds out completely when they are the indicated number of dB below GateOpen threshold.

Adjust **Floor** by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

Gate Thresholds

The two gate thresholds of the C1 allow you to adjust settings to minimize gate "chatter", the rapid opening and closing of the gate when signals are near the gate threshold.

GateOpen

Calibrated in dB below OdBFS. When in the *Gate mode*, **GateOpen** is the input threshold level above which the gate is opened to let the signal through. When in the *Expander mode,* **GateOpen** is the level at which soft knee expansion starts acting to a significant degree.

Adjust **GateOpen** threshold by dragging the upper equilateral triangle marker next to the Gate/Exp processor control level meter, or by dragging the right equilateral triangle marker under the input/output graph. You can also change **GateOpen** by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

GateClose

Calibrated in dB below OdBFS. In gate mode, **GateClose** controls the level below which gate is closed to stop signals from getting through. It is not used in expander mode. GateClose level is always equal to or less than GateOpen level.

Adjust **GateClose** by dragging the lower equilateral triangle marker next to the Gate/Exp control level meter, or by dragging the left equilateral triangle marker under the input/output graph. You can also change **GateClose** by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

> **Note:** In the expander mode of the Gate/Exp module, the **GateClose** and **Hold** controls have no effect, and the **GateOpen** control acts as a threshold control for the point where expansion action begins.

Attack

Calibrated in msec. Adjust the attack time of the expander/gate with time constants from 0.01 msec to 1 second. This adjustment is independent of the Attack time in the Comp/Exp module, although the two can be linked as described below. You can also change **Attack** by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

Release

Calibrated in msec. Adjust the release time of the expander/gate with time constants from 1 msec to 10 seconds. This adjustment is independent of the Release time in the Comp/Exp module, although the two can be linked as described below. You can also change **Release** by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrow keys.

Hold

Calibrated in msec. This control adjusts the length of time during which the gate is guaranteed to be held open.

Hold is particularly useful in creative gating applications in order to allow a fixed duration of a gated sound to be heard. It should be set at a low value unless you wish to guarantee a minimum duration of unprocessed transient to be let through. Please note that when the Gate/Exp module is in expander mode, **GateClose** and **Hold** controls have no effect.

EQ Mode switch

The EQ Mode switch is only available when sidechain is used and therefore is not seen in the "C1gate" component interface.

This has three modes: *Wideband, Sidechain* and *Split.*
In **Wideband** *mode*, the Gate/Exp operates as a conventional wideband gate or expander.

In **Sidechain** *mode*, the sidechain audio signal used to control the dynamics is first passed through the filters in the EQ module to make the dynamics respond only to a defined frequency band in the audio.

In **Split** *mode*, the sidechain audio signal is filtered, but the main signal is also split into two complementary bands: the *active band* (whose frequency response is shown in red on the graph) and the *passive band* (shown in blue on the graph). Only the active band is subjected to the Gate/Exp, while the passive band is not dynamically processed in any way. By this means, a dynamic equalizer, whose action is confined to the active band, is obtained.

When using the full C1comp/gate with sidechain, the EQ mode of the Comp/Exp module and the Gate/Exp module may be chosen independently of one another. For example, you can combine a wideband compressor with a high-frequency gate.

Gain reduction Meter

This shows the instantaneous gain reduction in dB (in red below 0 dB) or gain increase in dB (in orange above 0 dB) of the Gate/Exp dynamic processor.

Gate/expander control level meter

This shows the level of the signal controlling the processor (wideband or sidechain EQ). The meter time constants match the Gate/Exp attack and release times. As you adjust the GateOpen level, this meter shows the Peaks of sidechain signal.

Linking controls in the two dynamics modules

Certain controls of the two dynamics modules can be operated simultaneously, for ease of adjustment, by clicking the button between them.

These pairs of controls include:

- Threshold/GateOpen
- Attacks
- Releases
- EQ Mode

For more information on changing multiple controls at the same time, see the **WaveSystem Manual**.

Chapter 4 - The filter module

About the filter module

The filter module is used in the full "C1 comp/gate" and "+C1" components. Future components may be added with more sidechain control. Details will be in the Read Me file of any updates.

The C1 filter module splits the audio band into two complementary bands, **active** and **passive**. A graphical display shows the frequency response of the active band (in red) and the passive band (in blue), except when both processors are in wideband mode (when the filter module is not used).

Adjust the characteristics of the active and passive band by changing the filter controls: **Frequency** and **Q** (bandwidth). The active band may be any of four types of filters: bandpass, lowpass, high-pass, or bandreject.

These filters have rapid crossover rates. For example, bandpass or bandreject sidechain filters crossover at at least 18 dB per octave, and low and high pass at 36 dB per octave.

When used to bandsplit the audio signal, these filters are restricted to a Q of 0.6 or less to avoid the audible colorations associated with higher Q's. However, with a unique Waves phase compensation technology, the filters achieve good interband separation. They completely avoid frequency-modulation of the audio when used in bandsplit mode in order to avoid "out of tune" pitch shifting effects.

> **Note**: The frequency response graph is that of the more moderate Audio band filter. The SideChain filter in sidechain mode is always 3/2 times steeper than the audio filter in split mode.

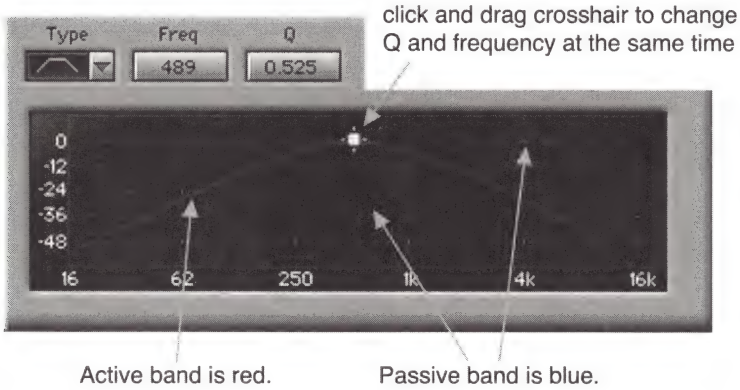

Filter Controls

Type

This selects whether the active band is *bandpass, band reject, high pass,* or *low pass,* and the button itself displays a schematic picture of the active filter response. The passive filter is of complementary type, passing bands which the active filter rejects and rejecting bands which the active filter passes.

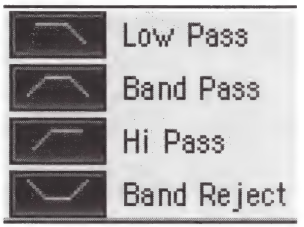

Frequency

Ranges from 16 Hz to 21 kHz. When *bandpass* or *bandreject* are selected, this is the center frequency of the respective pass or reject bands. When *low pass* or *high pass* types are selected, this is the crossover frequency between the two bands. Change the **Frequency** by clicking and dragging on the Value Window, by direct numeric entry, or by dragging the crosshair in the Filter graph to the left or right.

Q

Ranges from 0.1 to 0.6. In bandpass or bandreject filter types, **Q** controls the width of the band. The band is narrowest for highest Q. For low and high-pass filter types, **Q** determines the steepness of filter cut-offs. The cut-offs are steepest for the highest Q. Change the **Q** by clicking and dragging on the Value Window, by direct numeric entry, or by dragging the crosshair in the Filter graph up or down.

You can simultaneously control both **Q** and **Frequency** in the Filter graph by dragging the crosshair in the desired directions (left/right for Freq, up/down for Q).

> **Note:** The C1's **Q** calibration is the engineering textbook definition, which is different than the Q calibration found on normal audio parametric equalizers, such as the Waves Q10. Do not expect similar Q numbers for similar filtering in the C1 and Q10.

Chapter 5 - Other controls and I/O metering

Lookahead switch

About lookahead

Lookahead is not available in all C1 component plug-ins (which saves quite a bit of processing power). In this way, you use it only if you need it. It is included in any C1 component that has built-in sidechain support, that is to say, sidechain that is part of the plug-in itself, not part of the application, such as +C1, C1c/sc, etc.)

Lookahead processing is a method of delaying the audio signal with respect to control signals in order to minimize transient distortion caused by dynamic processing. It avoids premature cut-off of transients of gated signals, minimizes transient overshoot on compressed signals, and allows the use of longer attack times without unwanted side effects.

The *lookahead control* allows both dynamic processors to be operated with or without control signal lookahead. (The normal mode in most analog dynamic processors is without any control signal lookahead). **Lookahead** is a simple Yes/No switch with no other user controls. Lookahead time is automatically adjusted separately for the Comp/Exp and for the Gate/Exp module according to the attack times set up on each. At the 44.1 kHz audio sampling rate, the maximum lookahead time is approximately 7.5 milliseconds (the delay is exactly 340 samples, regardless of samping rate).

The use of lookahead generally means that longer attack times may be used without transient shape distortion becoming serious, even in gating applications. Such longer attack times result in lower nonlinear distortion and a cleaner sound. Note that lookahead processing does *not* affect the timing of signals from the start of a destructively processed sound file; therefore, you'd need to select at least a little before your audio to ensure the proper lookahead behavior. If the attack time is 14ms, extend your selection prior to the audio start by at least 14ms, etc.

Generally, you can use lookahead in all dynamic processing, unless you want the particular sound of modified transients obtained without lookahead. Sometimes, modified transients may be desired. For example, modified transients are effective when you want to gate an unwanted initial transient or to make the sound of sibilants bright when compressing speech.

But in compression applications, lookahead processing generally creates lower peak levels than nonlookahead processing, which makes lookahead processing the better choice when you want to maximize signal levels.

Limitations in available memory for lookahead mean that for longer attack times, especially above 15 or 20 msec, the lookahead mode is not as effective, although there is no reason not to use it.

Input/Output Graph

About the input/output graph

This graph shows the relationship between input and output levels for both the Comp/Exp and Gate/Exp modules. Both input and output scales are calibrated in dB from 0 to -100 dBFS. Input level is the horizontal scale, and output level is the vertical scale.

The input/output graph is particularly useful for visualizing the effect of the processing on signal dynamics, and for adjusting the processing. It includes metering of control levels for the Comp/Exp and Gate/Exp processor, displayed on the curves to aid in processor setup.

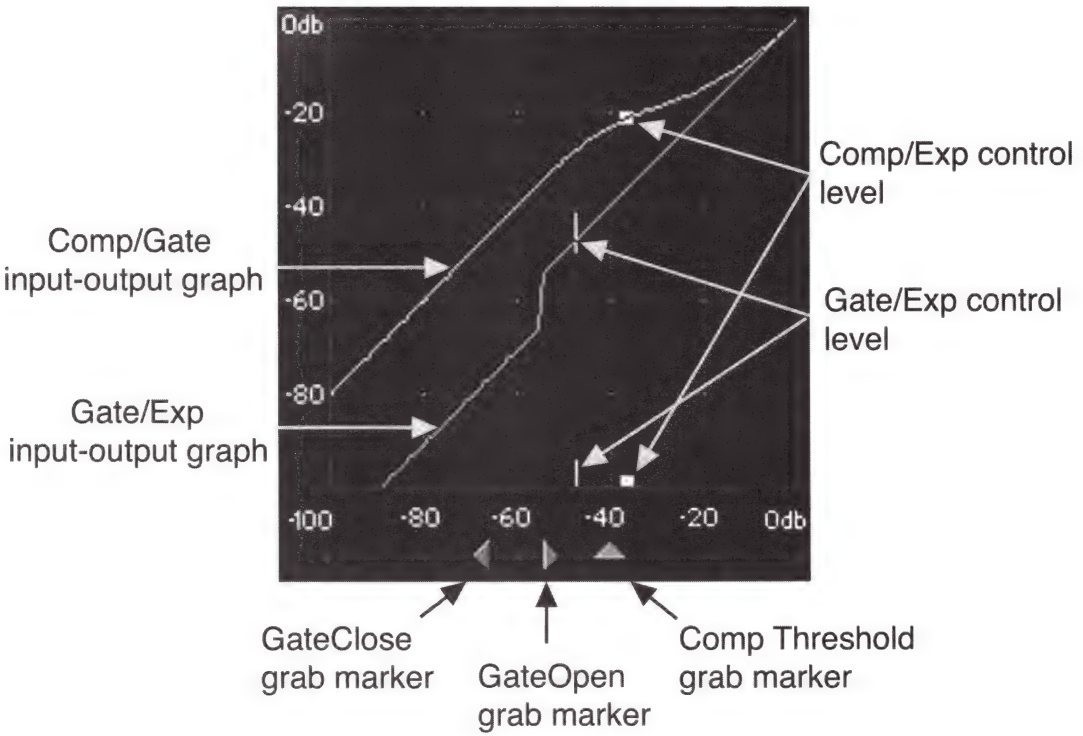

Comp/Exp display

Onscreen, the yellow curve shows Comp/Exp output level versus input level (depending on the EQ mode, this 'input level' will be the wideband or the sidechain level). Riding on the curve is a small square indicator showing the instantaneous level to which the Comp/Exp is responding. This provides an instant visual display of the actual effect on signal dynamics. To make adjustment of the threshold easier, a duplicate of the square moves along the lower edge of the I/O graph. Also on this lower edge is an upward-pointing yellow triangle. Drag it to adjust the threshold.

Gate/Exp display

The light blue curve shows the gate/expander processor output level versus input level (depending on the EQ mode of the processor, this 'input level' will be the wideband or the sidechain level). Riding on the curve is a small vertical line showing the instantaneous level to which the Gate/Exp is responding. This provides an instant visual display of the actual effect on signal dynamics. To make adjustment of the threshold easier, a duplicate of the line moves along the lower edge of the I/O graph.

Also on this lower edge, below the Gate/Exp's yellow threshold triangle, are two small blue triangles. Drag them to adjust the GateClose and GateOpen. As you drag the GateOpen, GateClose will follow it, maintaining the difference in dB you have set. For easy adjustment, click and drag GateOpen upward, pulling GateClose with it, until the gate opens the way you wish. Then, drag GateClose back down a few dB to create a hysteresis. This helps to avoid gate 'chatter'.

All three markers—Comp Threshold, GateOpen, and GateClose—may be adjusted up and down together by selecting all three and then dragging any one of them.

Keying mode

Stereo C1 components have keying modes, allowing one channel's energy to control the other's dynamics, or "key" it. Common applications are **ducking** and **gating**. **Ducking** reduces the gain of one signal when the other is above a threshold. It is normally used to make a background sound quieter when a foreground sound must be heard over the background. **Keyed gating** is essentially one signal controlling the gate of another. It is commonly used to gate a bass track with the kick, or to gate a continuous musical track with a percussion track in order to achieve a creative synchronized effect.

The Key button has three modes: Stereo (normal), L->R, and R->L. Click the button to cycle through the modes.

> **Note: In some platforms, to use the keying mode when processing destructively, you should select the channel you want Keyed, and that channel only.** In other words, do NOT select both Right and Left channels! When you Process in this mode, the channel you have selected will be Processed! If you have both channels selected, the 'keying' channel may be rewritten with **silence**. Therefore, if you still want to have the channel that is doing the Keying in the file after processing, make sure you have selected only the channel you want processed (the one being Keyed).

Stereo

Left keys Right

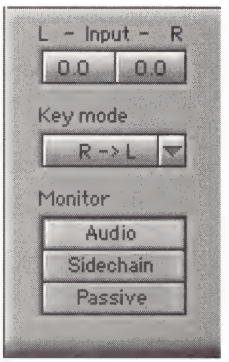

Right keys Left

Monitor switch

This allows audible monitoring of sidechain filtering and of the frequency processing. It controls what is heard at the output. Click the button to cycle through the three modes:

• Audio

Monitors the final output of the C1 dynamic processor and is the normal processing mode.

> ***Note: When Key Mode is in L->R or R->L mode, only the keyed channel is heard.***

• Sidechain

If either processor is in split or sidechain mode, this allows you to monitor the equalized sidechain signal. (When both processors are in wideband mode, the wideband audio sidechain signal is heard.) This is very useful in de-essing or noise filtering applications in order to more easily tune the sidechain filtering for optimal processing. Note: When Key Mode is in L->R or R->L mode, only the keyed channel is heard.

• Passive (unprocessed)

If either processor is in split mode, this allows you to monitor the sound of the passive (unprocessed) band, so that the effect of the processing is heard when the dynamic processor is at maximum gain reduction. This is useful in de-essing or noise filtering applications in order to listen to the frequency effect from the maximally processed signal.

Input and Output Levels

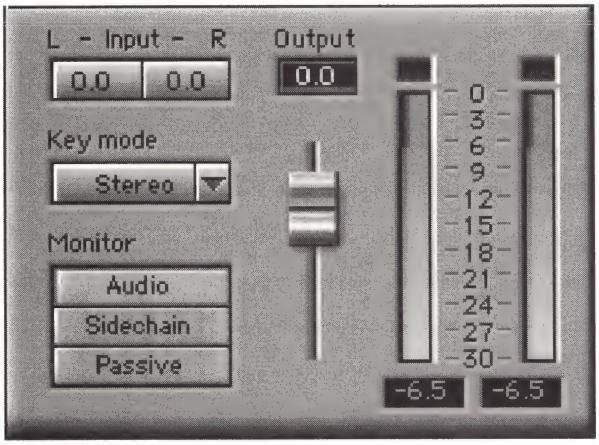

Adjust the input levels of either stereo channel by clicking the Input Value Windows and dragging with the mouse, entering values directly from the keyboard, or tweaking with the up/down arrows.

Adjust the output level by dragging the Output fader with the mouse, entering values directly from the keyboard, or tweaking with the up/down arrows.

Chapter 6 - Basics of operation

Use what you need

Select the C1 component for the job; this saves processing power and processing time. If you need a vocal compressor, then the C1comp is probably all you need. For a "hiss gate", you'll have to have the Gate with Sidechain, such as in the full C1 comp/gate or the +C1.

Not all possibilities will be discussed in this chapter. Some of the examples assume you're working with the full C1, but the controls are still the same, and most examples will apply to smaller components.

The Compressor module

To use only the basic compressor module, select the "C1comp" plug-in. If you're using the full C1 for this example, remember to set both **EQmode** switches to **Wideband** and the **Gate/Exp** bypass switch to **Bypass**. Make sure that the compressor module is not in Bypass mode.

Unless a special sound is desired, set the **Lookahead** switch to **Yes**.

Normally, set both input levels to maximum (0 dB), since this will never cause internal clipping problems in wideband mode. The input level adjustments are mainly useful for correcting channel imbalances (or intentionally creating them), since the output level fader can be used for overall level adjustments.

The C1 Comp/Exp module default settings are for basic compression. The ratio can be adjusted from 0.5:1 (expansion) to 50:1 (limiting). Beyond the 50:1 ratio are *negative ratios*, part of the unique *cancellation* mode of the Comp/Exp module.

Another distinctive aspect of this module is the **mid-level compression** that the compressor curve provides. To see this curve, set Ratio to about 4:1, and lower the Threshold to about -50dB. (See illustration on next page.)

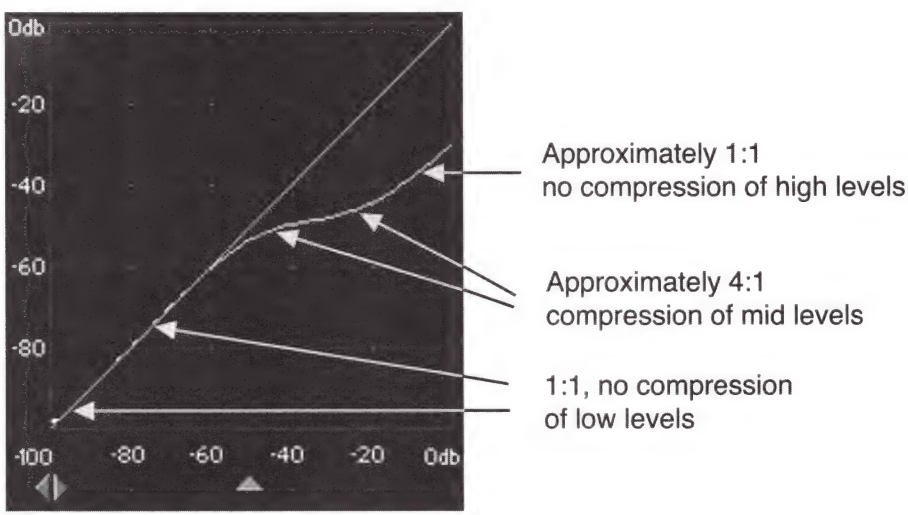

This **mid-level compression** allows you to increase loudness of a soundfile, or to make the sound more dense, *without overcompression of the peak dynamics.* Notice that the curve is not made of individual line segments but is a true soft-knee curve, providing smooth transition at all dynamic levels. This is ideal for post production, multimedia, and sweetening.

The Comp/Exp module is also capable of **negative ratios** (reverse compression). This occurs in the region beyond full limiting. Note how the level comes back down again over the threshold, and how, for more extreme reverse compression, it actually cancels out the signal completely at a level above threshold. This mode is also termed **cancellation mode** because of this property. It is useful if you wish to bring down high level input signals even further than would with a limiter.

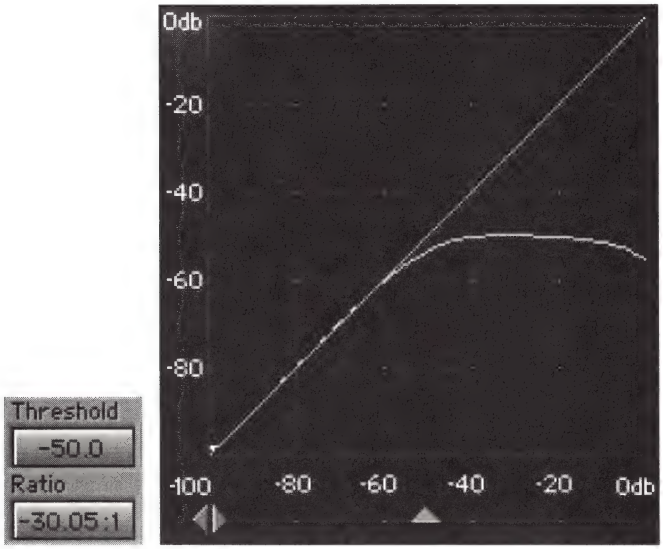

Cancellation mode

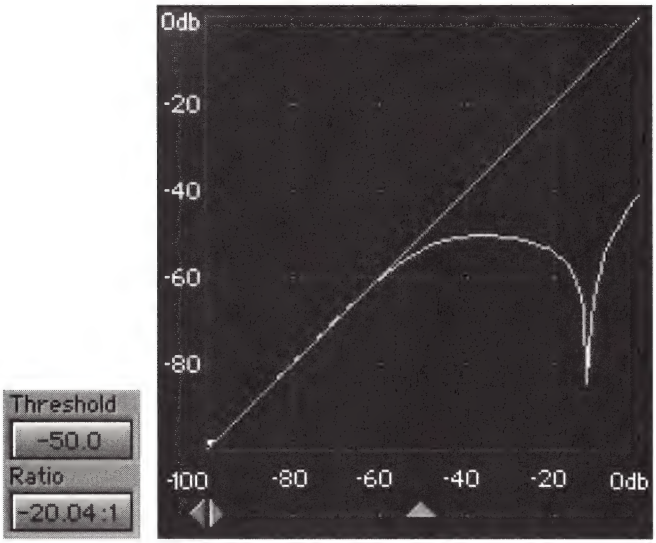

Reverse compression

LowRef/PeakRef mode

There are two different ways of setting up the Comp/Exp output gain structure, using respectively LowRef and PeakRef adjustment modes. The adjustment mode may be selected with the LowRef/PeakRef reference level switch.

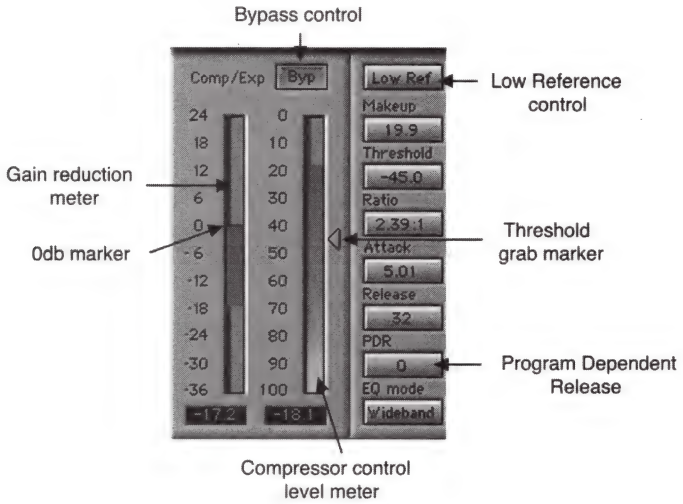

Which level reference mode should you use?

It depends on what you want to do with the compressor. Waves has provided the two modes to make it easier to set up for different applications. While no absolute rules can be given, and everyone has their own preferred ways of working, the following is a rough guide.

LowRef mode is the classic compressor mode with which most people are familiar. Since this involves no relearning of standard methods of operation, this may be preferred by many users experienced in using audio compression.

LowRef mode is typically used if the threshold is not too "deep" (that is, if a majority of the signal dynamics is well below the threshold setting). It is also recommended if large compression ratios such as 4:1 and higher are being used. This is especially the case when the compressor is in split mode.

PeakRef mode would typically be used if you have very low threshold settings (below -30 dB). This will be the case if you wish to compress (or expand) mainly the lower-middle or middle levels of the input dynamic range, leaving both very low and high levels uncompressed (i.e. "mid-level" compression applications). Typically in these applications, the ratio will lie between about 0.65:1 (expansion) and 3:1 (compression).
Usually, in these applications, the threshold setting will be low, somewhere in the middle or low-middle of the signal's dynamic range.
You can see the relative setting of threshold and input signal level by comparing the threshold marker with the control level meter indication, or by comparing the threshold marker with the level indicated by the square indicator on the I/O graph.

Adjustment of the compressor in LowRef mode

This is the traditional mode for a dynamics processor. When the Comp/Exp module is used as a compressor, as you lower the threshold, the output goes down. Adjust the Makeup gain value to compensate for the gain reduction.

Adjustment of the compressor in PeakRef mode

In PeakRef mode, first adjust the Makeup gain value so that it reads 0 dB or slightly less, such as -3dB. This ensures that output peak levels are roughly the same as input peaks.

When the Comp/Exp module is used as a compressor, as you lower the threshold, the output gain will increase to approximately compensate for the gain reduction.

After compressor adjustment, some additional trimming of the Makeup gain may be needed to prevent output peak clipping. (The Makeup gain controls the output of the Comp/Exp module.)

How to adjust PDR

In general, fixed release times are best for a somewhat small range of dynamic action. In reality, a release time that automatically adjusts to the sound material is sometimes desirable. You will find that the *Program Dependent Release* feature of the C1 allows for a controllable automatic release time.

When **PDR** is at its lowest value (0), it is disabled, and the **Release** time of the Comp/Exp module is fixed.

Set the **Release** time to the longest time desired, then set **PDR** for the compressor action that gives the desired effect.

The actual Comp/Exp release time is shortened whenever a transient occurs that is shorter than the **PDR value** (calibrated in millseconds). As the value of PDR gets larger, it takes longer sustained energy to get PDR to work. For example, if your soundfile has a sudden sustained, high-level area followed quickly by a very quiet area, then you would set PDR longer to avoid 'hole-punching'. Of course, if you want to have creative 'pumping', now you can control it with careful adjustment of PDR.

If you are using a long release time but have sounds with short transients (typically, a few milliseconds) such as speech sibilants and 'plosives' or drumstick attack sounds, the short transients will reduce gain for a long time without PDR. You can set the PDR just long enough (perhaps 10 or 20 msec) to allow rapid recovery of gain after these to prevent them from reducing level too much.

In applications where very long release times are used, you may wish to use PDR times as long as perhaps 300 msec so that the release time is short for all individual sounds having a limited duration of up to a fraction of a second, with the fixed **Release** time acting only after sustained periods of high-level sound. As an example, this would help to prevent reverberation from being over-exaggerated by the use of compression.

You will find that on recordings in very reverberant acoustics (e.g. classical music), PDR can allow significant compression without the reverb tails being brought up too rapidly.

Chapter 7 - The Gate/Expander module

To use only the basic Gate/Exp module, select the "C1gate" plug-in. If you're using the full C1 for this example, set both **EQmode** switches to **Wideband** and the **Compressor** bypass switch to **Bypass**. Make sure that the Gate/Exp module is not in Bypass mode.

For best results, unless a special sound is required, the **Lookahead** switch should be set to **Yes**.

Normally, set both input levels to maximum (0 dB), since this will never cause internal clipping problems in Wideband mode. The input level adjustments are mainly useful for correcting channel imbalances (or intentionally creating them), since the output level fader can be used for overall level adjustments.

To use as a Wideband Gate

Set the **Gate/Exp** mode to *Gate.*
Set **EQ mode** to *Wideband.*
Set **Floor** to *-infinity.*
Then drag the **GateOpen** threshold up to a moderate level so that the Gate opens the way you want.
Set the **Attack** and **Release** times for the desired gate action.

To minimize gate "chatter" (or rapid opening and closing of the gate when the signal is near **GateClose**), lower the **GateClose** value a few dB below the GateOpen value, (for instance, -3dB lower). This will close the gate at a slightly lower level than the level necessary to open the gate, allowing the sound to die out a bit more before closing.

In the following illustration, the **GateOpen** and **GateClose** are both set to -41.7dB. Note that the Gate/Exp line in the Transfer Graph shows the Gate clearly, and all signals below -41.7 are gated.

To set the **Hold** time, adjust the Hold value (calibrated in mS) by clicking and dragging on the Value Window, by direct numeric entry, or with the up/down arrows. The Gate will stay open for this time period before the Release time occurs. In other words, once the level goes above the GateOpen threshold, the gate will remain open for the Hold time specified, then the Release time will begin. The Hold time is essentially a 'delay before release', such as the "S" part of an "ADSR" device.

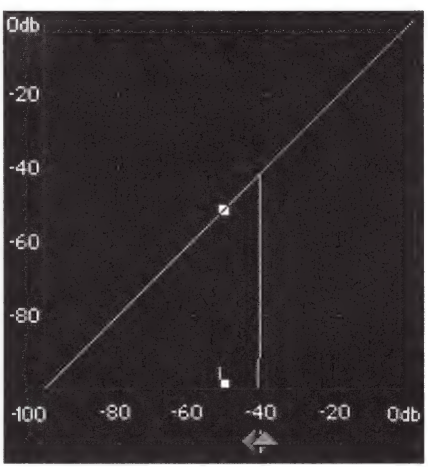

To use as a Wideband Expander

Simply set the **Gate/Exp** mode to *Expander* and leave all other settings the same. The only difference is that now the module is a *downward expander*, with a ratio of 1:2. A downward expander will give you a more subtle, less abrupt alteration of the dynamics as the signal level drops.

In the following illustration, note the sloping Gate/Exp line showing the 1:2 expansion.

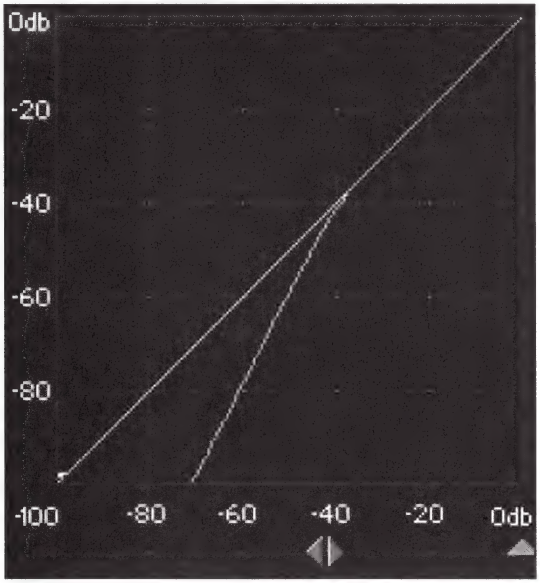

To use the Floor control

Floor (calibrated in dB) controls the lowest level a sound will be taken to when Gated or Expanded. The default is (-infinity), which means that the level can be taken down to digital silence. However, many times it is desirable to leave a little background noise "room tone" instead of completely gating all sound. This is when you can use the **Floor** control.

To adjust **Floor**, just click and drag on the Value Window, enter numeric values directly, or use the up/down arrows.

In gate mode, the **Floor** control ranges from -infinity to +12dB. To retain a little "room tone", set the Floor value to about -30, as in the following illustration. It is also possible to *increase* the level below the GateOpen threshold by setting Floor to positive values. This will be discussed in the chapter **Putting it all together.**

In expander mode, the **Floor** control also goes beyond the *-infinity* setting to *negative polarity settings* (-10 N to -100 N). This gives the expander a steeper expansion ratio below threshold. The expander will cancel sounds out completely when they are the indicated number of dB below GateOpen threshold.

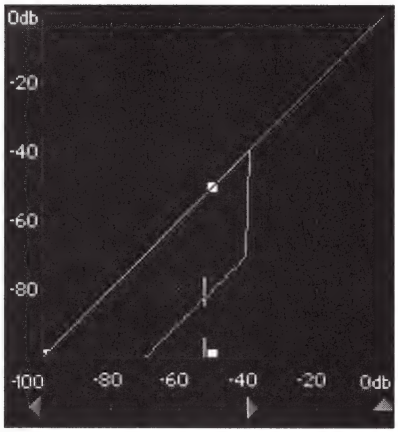

The Filter module

The great power of the C1 comes from its **Filter** module, which can be used for Sidechain EQ or the unique *bandsplit modes.* For example, this bandsplit configuration allows you to compress just the bass while gating wideband, or compress wideband while de-essing. These are just two of many possible *combination setups.* The C1 Setup Library presents a wide array of dynamic tools based on this *bandsplit mode.* You will most likely find this library and its documentation very usable. We strongly suggest you take the time to explore it.

First, the basics about the filter/equalization module. The Filter has bandpass, bandreject, highpass and lowpass filter types. The Filter's Graphic Display shows the filter and status of the active/passive bands. In the following illustration, the Filter is in *Bandpass* mode, with a center **Freq**uency of 489Hz and a **Q** (bandwidth) of 0.525.

Click on the **Type** control to cycle through the 4 filter types. The **Freq** and **Q** controls can be adjusted by the usual methods: clicking and dragging on the Value Window, entering values numerically, or using the up/down arrow keys.

> **Note:** You may also adjust the **Q** and **Freq** simultaneously by dragging the crosshair, as shown in the illustration. Move the mouse vertically to adjust the **Q**, horizontally to adjust the **Freq**.

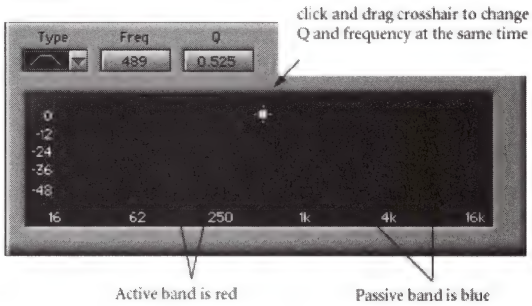

Chapter 8 - Sidechain Tutorial

To really start putting it together, first we'll present a classic wideband compressor and gate. Here are the steps for this basic setup:

Select the full C1 comp/gate with sidechain (you can use +C1).
Set both modules' **EQ mode** to *Wideband.*
Set both modules so they are not Bypassed.
Set the desired **Ratio**.
Drag the Comp/Exp **Threshold** to obtain the desired gain reduction.
Now drag the Gate/Exp **GateOpen** to obtain the desired gating action.
Adjust the **Attacks** and **Releases** to appropriate settings for the file.

It's this simple.

Next step

Now we'll bring the **Filter** into the configuration.

First, here's how a single dynamic module and the **Filter** combine for sidechain. This is the traditional sidechain compression technique used for de-essing.

- Put the Gate/Exp in *Bypass.*
- Set the compressor for traditional compression *(ratio 2:1, attack 5ms, release 30ms)*
- Set the Comp/Exp **EQ mode** to *Sidechain.*

Now the Comp/Exp module will compress according to the amount of energy in the filtered signal. You can adjust the filter to 'tune in' on the level of energy to which you want to make the compressor sensitive.
Listen to the sidechain filter output by clicking the Monitor button once so that it says *Sidechain.*

Watch the blue Comp/Exp Control Level VU meter. Adjust the Threshold grab marker (the triangle) and drag it so that the VU meter reading moves above and below the marker. The red Comp/Exp Gain Reduction meter will show gain reduction whenever the Control Level goes above the Threshold. With this interface, you can set the sidechain filter and adjust the threshold very quickly.

Now set the **Monitor** mode back to *Audio* by clicking the button to hear the de-essing process.

Better (bandsplit) de-esser

Second, here's another simple de-esser, but this time in *bandsplit mode.* The settings are all exactly the same, but instead, you will compress only the 'ess' part of the band, instead of the entire wideband signal.

- Put the Comp/Exp module in *Split mode* by clicking the **Eq mode** button until *Split* is shown on the button.
- Now adjust the module for desired de-essing action.

This is not necessarily the optimum de-esser! It's just an example. Check the "Better De-esser" Factory Preset.

This Preset can be easily made into a **De-Popper** by changing the sidechain Freq to 120!

You can combine the **Filter** with the Gate/Expander module in the same way. If you have a soundfile that has lots of low end energy that keeps triggering the Gate at the wrong times, use the **Filter** in the *Highpass* mode to help prevent these low frequencies from triggering the Gate.

An ideal application is to have a sidechain gate to clean up drum tracks. Use a Highpass filter to isolate the snare, a Bandpass filter for toms, and lowpass to gate the kick.

Chapter 9 - Intro to Combination setups

Combination setups are for use with both modules plus the Filter module.

Using both dynamic modules with one module in Sidechain or Split mode offer the most flexible possibilities. There is only one Filter, but the two dynamic modules can perform separate processing while using the Filter in very different ways.

For instance, with the Filter in Highpass mode, you can have the Comp/Exp module in sidechain mode as a de-esser, and have the Gate/Exp module in split mode as a de-hisser. Both modules would be controlled by the energy measured in the highpass band. The difference is that the compressor would be acting on the wideband signal, the gate only affecting the highpass band.

Combining the three modules of the C1 will give you considerable flexibility—in fact, so much flexibility that it is fastest to start with the Setup Library for many applications, then tweak the settings for your work. To introduce you to this complex aspect of C1, the examples in the following section use setups from the C1 Setup Library.

Classic Combinations

High-level Compressor + De-Hisser

The first two combination examples were for de-essing. Now, here's the third, but this combination is designed as a "de-hisser". It uses both dynamic modules and the Filter. This time, the Gate/Expander will be the de-esser, while the compressor will remain a classic wideband compressor. In this setup, the real essence of the C1 is demonstrated: two independent effects which may be set up, optimized, and controlled independently or together.

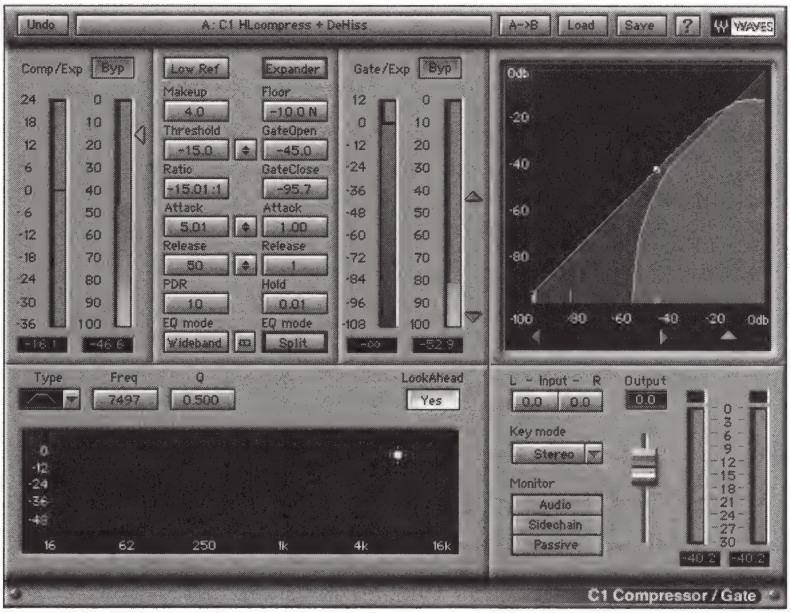

The Comp/Exp module is set up as a wideband compressor. Note that the compression Ratio is set to a negative value to bring down peaks a little further. This is very effective when any cancellation nulls lie beyond 0 dB. Of course, you can have traditional compression with a more traditional Ratio (for example, 3:1).

Note that the Makeup gain is +4.0dB. However, the use of compression with this Makeup gain will bring up any background hiss. The **Gate/Exp** module is set up as a de-hisser to minimize this background hiss with the least effect on the sound. It is a low-level expander whose action is confined to a carefully-tuned frequency band centered around 7 kHz.

In this setup, the **GateOpen** threshold setting is the critical adjustment for de-hissing. Adjust it up and down until you find the highest setting that has no significant unwanted effect on tonal quality. This is the optimum de-hiss setting. The setting is a trade-off between loss of treble and hiss reduction. On most speech and popular music, there is usually a good setting that has a subjective hiss reduction of 6 dB or more but with little tonal effect . The filters chosen here for de-hissing are generally suitable for hiss with a white, pink or blue spectrum, requiring only perhaps a kHz up or down tuning of frequency. The tonal effect of maximum hiss reduction can be heard by setting the **Monitor** to *passive* mode. The **Q** setting turns out to be fairly critical — lower Q (wider bandwidth) will filter out more hiss — but modulation noise (variations of the hiss up and down with the wanted signal) will become more audible. A **Q** of 0.5 is generally the best compromise, but up to 0.6 can be used. The optimum filter settings usually end up close to those in the setup.

Main controls for Compressor + De-Hisser:
Comp/Exp *set to high-level compression.*

- **Makeup** as needed
- **Threshold** -20 to 0
- **Ratio** -50 to -5 as needed
- **Attack**, **Release**, **PDR**: As needed

Gate/Exp *module set to de-hissing*

- **GateOpen** threshold. -25 to -80. Set by ear to leave signal treble balance unaltered while diminishing hiss.

Other controls:

- **Frequency**: In range 5500 to 9000 Hz for lowest perceived hiss
- **Q**: In range 0.45 to 0.6 for best trade-off of modulation noise versus hiss reduction
- **Output Level**: (has same function in this setup as makeup)

Compress + De-reverb

From the Setup Library, load the file **Compress+De-reverb**. Try this with a voiceover with too much room reverberation. This setup uses the compressor for mid-level wideband compression in order to increase loudness and density of a voice. However, this brings up lower level sounds, including reverberation. For post and multimedia applications, setting the Gate/Exp module for *splitband downward expansion* enables you to remove some of this reverb.

This means that you can reduce the volume of the lower frequencies — the ones responsible for the 'cloudy' sound of compressed reverb — by expanding their dynamics downward. Note that the **Filter** is set for a crossover frequency of 400Hz.

Main controls for Compress+De-reverb:

Comp/Exp *set to mid-level compression.*
- **Makeup as needed**
- **Threshold -70 to -20.**
- **Ratio 1:1 to 3:1**
- **Attack, Release, PDR: As needed**

Gate/Exp *set to bass de-reverberation*
- **GateOpen threshold:** -40 to -15 set by ear to leave signal bass balance unaltered while diminishing reverb bass
- **Attack, Release, Hold:** As needed - usually Attack not too short (to avoid distortion of the bass frequencies) and Release and Hold not too long (to allow quick response to low level in the bass)

Other controls:
- **Frequency:** In range 200 to 600 Hz for best sound
- **Output Level** (has same function in this setup as makeup)
- **Lookahead**

GateOpen/GateClose

By dragging GateOpen, you may vary the two settings up and down together with a constant dB difference between them. You can also adjust them together by dragging the upward-pointing grab marker (GateOpen) on the Gate/Exp control level meter, **or** by dragging the right-pointing arrow under the I/O graph (see illustration on facing page).

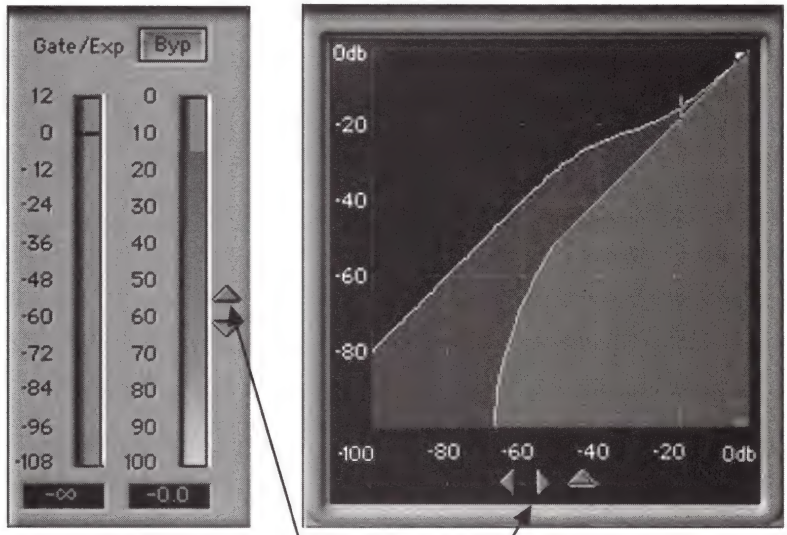

Drag either of these grab markers (GateOpen) to change GateOpen and GateClose at the same time

Please see the WaveSystem Manual for complete details of these and other control features.

Chapter 10 - Setup Library listing

Simple Setups

C1 Multimedia Speech 1
C1 Multimedia Speech 2
C1 EQ + Gate
C1 EQ + LLexpander
C1 EQ + LLcompressor
C1 Compressor + EQ

Compression

C1 Speech Compress/Expand 1
C1 Compressor + De-Reverb
C1 HLcompress + DeHiss

Noise Reduction

C1 Noise Reducer
C1 Noise Reducer 2
C1 Gate Noise Reducer
C1 Rumble Reducer
C1 High-pass filter

De-Essers

C1 Compressor + DeEsser 1
C1 De-Ess + De-Hiss 1
C1 De-Ess + LLcompress
C1 De-Ess + EQ + De-Hiss 2
C1 De-Hiss + EQ + De-Ess 3
C1 De-Hiss + EQ + De-Ess 4

Enhancers

C1 Speech Enhancer
C1 Treble Enhancer
C1 Bass Enhancer
C1 Bass/Treble Enhancer
C1 CheapSystem Enhancer

Keying Setups

C1 Ducking 1
C1 Ducking 2
C1 Ducked EQ 1
C1 Ducked EQ 2
C1 Keyed Gate
C1 Keyed Expander
C1 Keyed EQ Expander

Chapter 11 - Conceptual block diagrams of all EQ mode options

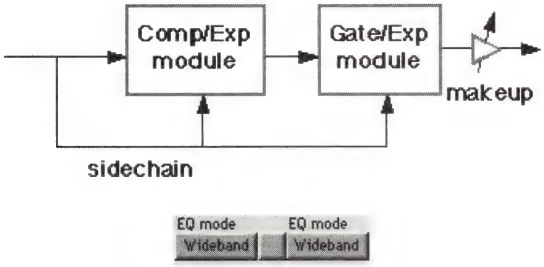

Both modules in Wideband mode.

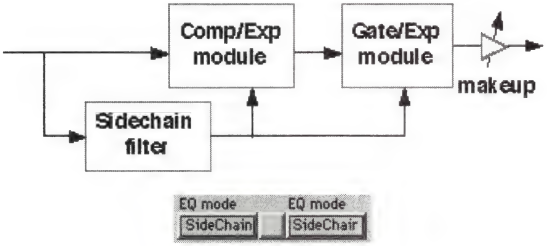

Both modules in Sidechain mode.

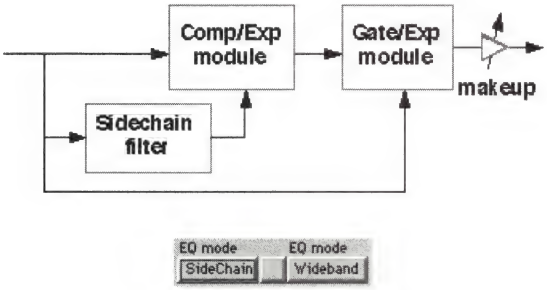

Comp/Exp module in Sidechain mode, Gate/Exp in Wideband.

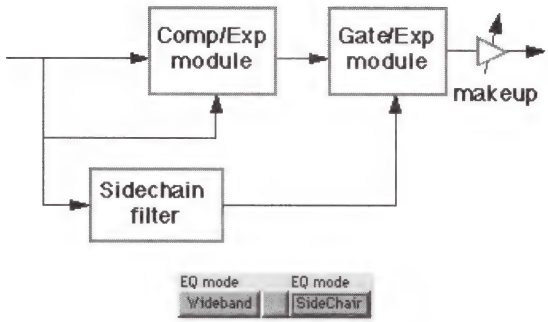

Gate/Exp module in sidechain mode, Comp/Exp module wideband.

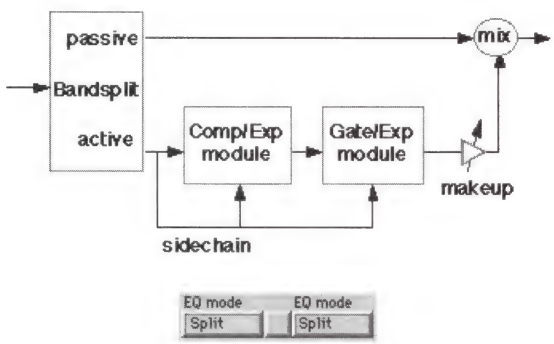

Both modules in bandsplit mode.

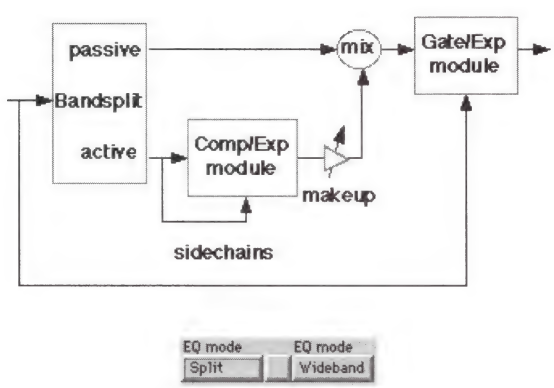

Bandsplit Comp/Exp module, with wideband Gate/Exp module.

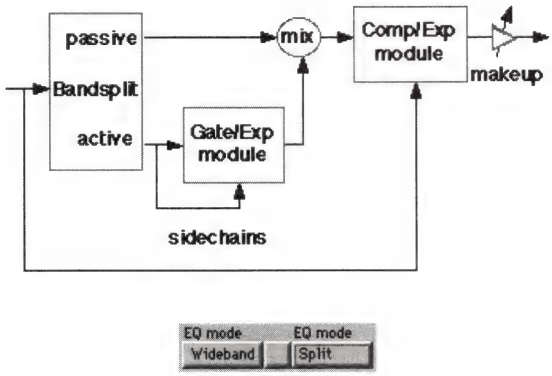

Bandsplit Gate/Exp module with wideband Comp/Exp module.

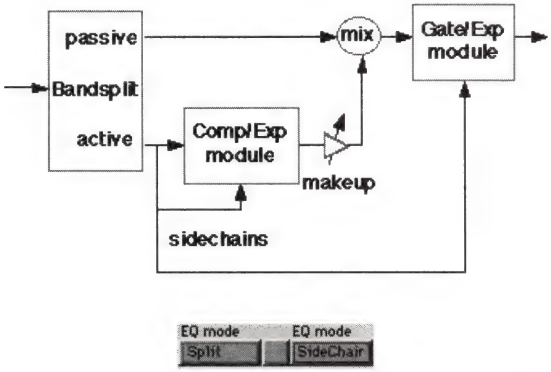

Comp/Exp module in bandsplit mode, Gate/Exp module in sidechain mode.

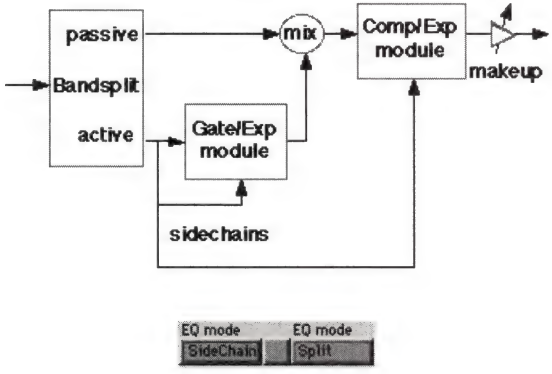

Gate/Exp module in bandsplit mode, Comp/Exp module in sidechain mode.

Key modes

The following show "R keys L" modes. The "L keys R" modes simply interchange the L and the R inputs.

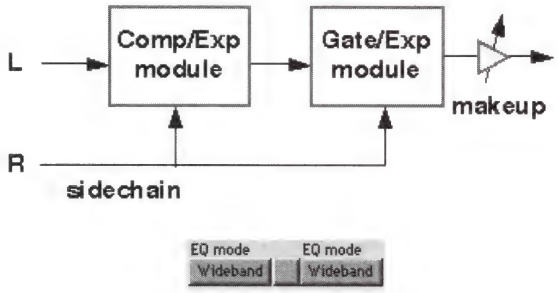

Both processors in wideband mode. R keys L.

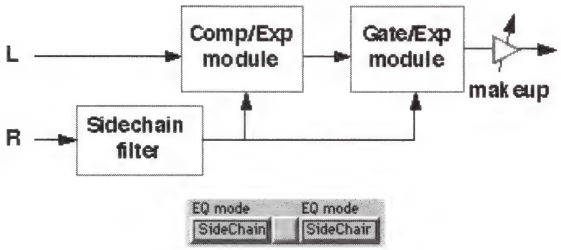

Both processors in sidechain mode. R keys L.

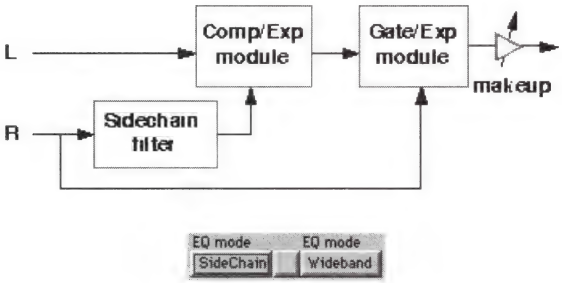

Comp/Exp module in sidechain mode.
Gate/Exp module in wideband mode. R keys L.

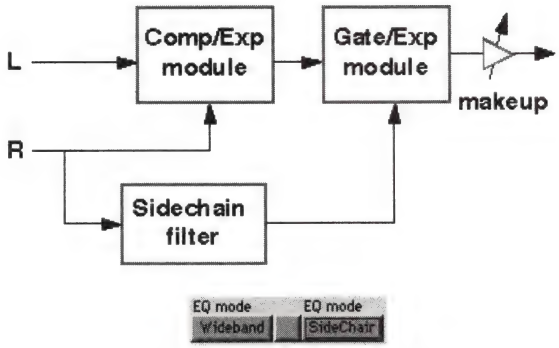

Comp/Exp module in wideband mode.
Gate/Exp module in sidechain mode. R keys L.

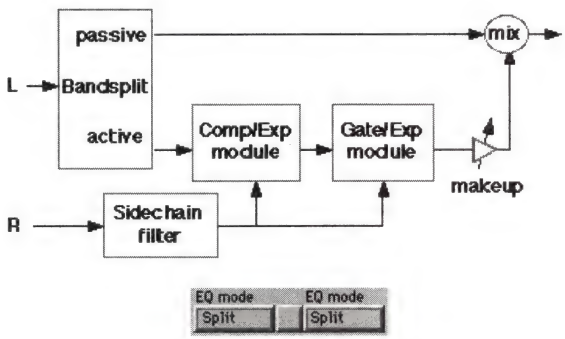

Comp/Exp module and Gate/Exp module both in Bandsplit mode. R keys L.

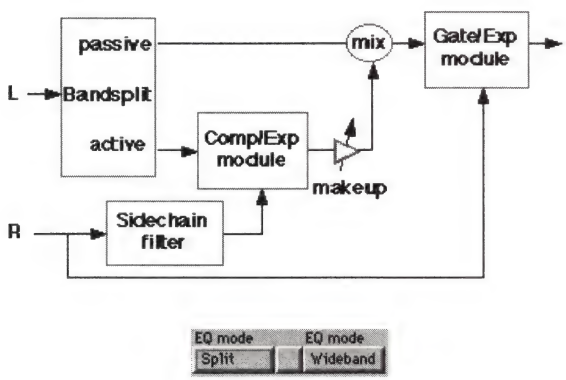

Comp/Exp module in Bandsplit mode.
Gate/Exp module in wideband mode. R keys L.

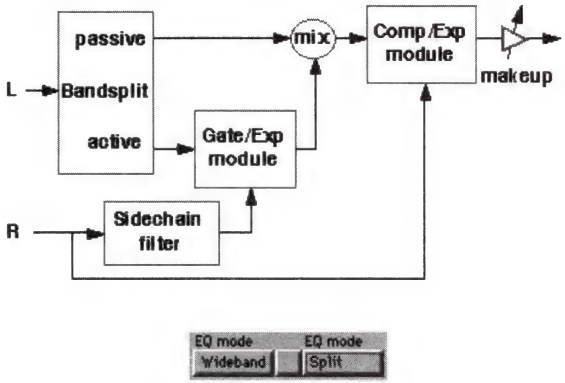

Gate/Exp module in Bandsplit mode. Comp/Exp module in wideband. R keys L.

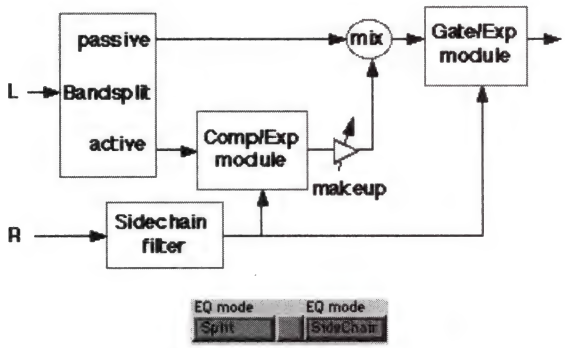

Comp/Exp module in Bandsplit mode. Gate/Exp module in sidechain. R keys L.

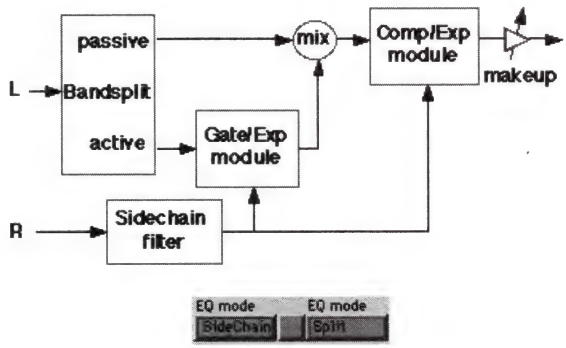

Gate/Exp module in Bandsplit mode. Comp/Exp module in sidechain. R keys L.

Chapter 12 - C1 Technical Specifications

Digital operating resolution.

Internal computations: 24 bits. Bandsplit filters incorporate internal noise shaping. Control signals are computed to 24 bit accuracy every sample to eliminate zipper and quantization noise.
Input signal: 16 or 24 bits.
Output signal: 24 bits fixed (TDM/DSP systems), 32bit float (native)

Digital Resolution enhancement: IDR normal type 1 noise shaped dithered quantization, at 16 or 20 bits, plus unmodified 24 bit output.

Primary sampling rates: 44.1 or 48 kHz, but supports other rates. C1comp and C1gate (without sidechain and lookahead) now support 88.2 and 96kHz rates in native applications.

Compressor/Expander module

Type: Soft knee with gain reduction 3 dB at threshold for limiting mode. Compression occurs over about 30 dB range above threshold, then ratio reverts to 1:1.
Ratio: 0.5 expansion to full limiting plus negative ratio settings.
Threshold: 0dBFS (full scale) to -100 dB in 0.1 dB steps.
Attack time: 0.01 to 1000 msec logarithmic scaling.
Release time: 1 to 10000 msec logarithmic scaling.
PDR time (length of transients for which release time is shortened): 1 to 1000 msec.
Makeup gain: -40 dB to + 40 dB.
Gain reference modes: Low-level referenced or peak-level referenced.
Lookahead: optional, linked to attack time.
EQ operating modes: Wideband, Sidechain EQ, Bandsplit mode.
Compressor Bypass switch: yes.
Metering: Sidechain input levels (100 dB range), inc. maximum reading, Gain reduction/increase, inc. maximum reading.
Graphical displays: Input/Output curve over 100 dB, with on-curve level metering. Sidechain/bandsplit frequency responses.

Gate/Expander Processor

Operational modes: Gate or 2:1 soft-knee expander.
Floor gain level: -infinity dB to +12 dB and (for expander only) negative polarity floor gains from -10N to -100N dB.
Expander/GateOpen threshold: 0dBFS (full scale) to -100 dB in 0.1 dB steps.
GateClose threshold: 0dBFS (full scale) to -100 dB in 0.1 dB steps with time-related hysteresis (gate only).
Attack time: 0.01 to 1000 msec logarithmic scaling.
Release time: 1 to 10000 msec logarithmic scaling.
Hold time: 0.01 to 5000 msec logarithmic scaling (gate only).
Lookahead: optional, linked to attack time.

EQ operating modes: Wideband, Sidechain EQ, Bandsplit mode.
Gate/Exp processing Bypass switch: yes.
Metering: Sidechain input levels (100 dB range), inc, maximum reading, Gain reduction/increase, inc. maximum reading.
Graphical displays: Input/Output curve over 100 dB, with on-curve level metering. Sidechain/bandsplit frequency responses.

EQ Section

Filter types: Bandpass, bandreject, High-pass, Low-pass
Q: 0.1 to 0.6.
Frequency: 16 Hz to 21 kHz logarithmically scaled.
Bandsplit filter type: IIR phase compensated bandsplit internally noise shaped.
Sidechain filter type: 6th order IIR.
Graphical display: Sidechain/bandsplit frequency responses.
Audio Monitoring: Processed sound, sidechain signal, "passive band" signal.

Other

Lookahead processing latency (only in components with the lookahead button): 340 samples (7.7 msec at 44.1 kHz sampling rate, 7.08 msec at 48). Processed files in native applications (non-DSP) may show latency.
Keying modes: Ordinary linked stereo, Left keys right, Right keys left.
Output metering: stereo peak-reading bar meters 0 to -30 dB, also indicating maximum level and number of clips (in +C1).
Input level: Separate left and right gains 0 to -24 dB in 0.1 dB steps.
Output level: slider gain +18 to -40 dB in 0.1 dB steps.
Overall bypass mode: Yes, with gain adjustment +12 to -24 dB for A/B comparisons.
Number of setups available: Currently loaded: Two. In memory: unlimited.
Real-time preview mode at 44.1 and 48 kHz sampling rates.

Features and specifications subject to change without notice.

C1
Setup Guide

Table of Contents

Chapter 1 ..... Introduction 2

Chapter 2 ..... Basic Use of Setups 7

Chapter 3 ..... Basic Dynamic Processors .. 15

Chapter 4 ..... Simple Setups .. 21

Chapter 5 ..... Compression .. 26

Chapter 6 ..... C1 Noise Reduction .. 31

Chapter 7 ..... DeEssers .. 38

Chapter 8 ..... Enhancers .. 43

Chapter 9 ..... Keying Setups .. 50

Chapter 1 - Introduction

The Setups Library's purpose

The C1 at heart is simply a compressor and expander/gate with some added filters, but its unique design, evolved through many years of experience and refinement, makes it a versatile tool for any number of specialized and useful audio processing tasks. Moreover, its performance at most of these tasks is state-of-the-art, either comparable to or superior to dedicated professional products available elsewhere. Many of the tasks it can do are unique to the C1.

The C1 Setups Library allows the C1 to be configured to become any of a number of dedicated processors. The Library consists of a variety of setup files, each of which when loaded turn the C1 into a specialized problem solver or creative tool.

What can you do with the C1?

Among the tasks the C1 can perform using the setups library are the following, with relevant setups and their chapters listed:

- Conventional high-level compression

Chapter 3:	C1 Classic Compressor/Gate
	C1 Classic Compressor/Expander
Chapter 5:	C1 HLcompress + DeHiss
Chapter 8:	C1 HLcompress+bass/treb enhance
	C1 HLcompress+treble enhance

- Conventional Low level expansion

Chapter 3:	C1 Classic Compressor/Expander
Chapter 4:	C1 EQ + LLexpander
	C1 Multimedia Speech 2

- Versatile conventional gating

Chapter 3:	C1 Classic Compressor + Gate
	C1 Classic DeEsser + Gate
Chapter 4:	C1 Multimedia Speech 1

- De-esser - reducing excessive speech or vocal sibilance

Chapter 3:	C1 Classic DeEsser + Gate
Chapter 7	- all setups

- Noise reducers - reducing background noises

Chapter 5:	C1 HLcompress + DeHiss
Chapter 6	- all setups
Chapter 7:	C1 De-Ess + De-Hiss 1
	C1 De-Ess + EQ + De-Hiss 2
	C1 De-Hiss + EQ + De-Ess 3
	C1 De-Hiss + EQ + De-Ess 4

- De-reverberation - reducing the audible effect of room reverberation, especially on speech

Chapter 5:	C1 Compressor + De-Reverb

- Hiss removal

See Noise Reducers above

- Rumble removal

Chapter 6:	C1 Rumble Reducer
	C1 High-pass filter

- Mid-level compressor - making sounds louder without compressing the dynamics of loud sounds

Chapter 4:	C1 Multimedia Speech 1
	C1 Multimedia Speech 2
Chapter 5:	C1 Compressor + De-Reverb
Chapter 8:	C1 MLcompress+bass/treb enhance
	C1 MLcompress+treble enhance

- Low level detail enhancer - bringing up the level of quiet sounds to make them more clearly audible

Chapter 4:	C1 EQ + LLcompressor
Chapter 7:	C1 De-Ess + LLcompress
Chapter 8:	all setups

- Spectral enhancers - giving sounds more depth, brightness or impact without the unpleasant side effects of conventional equalization or the harshness and artificiality of some commercial enhancers.

Chapter 8:	all setups

- Dynamic equalization - permitting a sound to be equalized in different ways at different sound levels, giving effects impossible with a simple equalizer.

Chapter 8:	all setups

- Speech/Vocal intelligibility enhancer - makes speech and vocals more clearly audible under difficult listening conditions without sounding artificial under ideal conditions

  Chapter 8: C1 Speech Enhancer

- Louder multimedia files - allows optimum sound quality to be obtained even from 8 bit 22 kHz multimedia files without "squeezing the life out of" the original sound file.

  Chapter 4: C1 Multimedia Speech 1
  C1 Multimedia Speech 2

- Ducking of one signal keyed by another

  Chapter 9: C1 Ducking 1
  C1 Ducking 2

- Gating of one signal keyed by another

  Chapter 9: C1 Keyed Gate

- Expansion of one signal keyed by another - ideal for creating convincing rhythmic tracks out of any continuous sound effect.

  Chapter 9: C1 Keyed Expander

- Duck-EQ - keyed ducking within a frequency band of one signal keyed by another, to allow space to be created in a mix when several sounds conflict in the same frequency band.

  Chapter 9: C1 Ducked EQ 1
  C1 Ducked EQ 2

- Creative keyed equalization of one signal by another, allowing the dynamics of a key signal to alter the sound of a second signal

  Chapter 9: C1 Keyed EQ Expander

Combinations of effects

Also, the C1 allows one-pass real-time preview and processing of many combinations of two or even three effects at the same time, for example allowing you to compress a sound file and at the same time to remove side-effects such as increased audibility of noise or room reverberation. These combined effects include:

- **Compressor + Low-level expander**
  Chapter 3: C1 Classic Compressor + Expander
  Chapter 4: C1 Multimedia Speech 2

- **Compressor + gate**
  Chapter 3: C1 Classic Compressor + Gate
  Chapter 4: C1 Multimedia Speech 1

- **Compressor + equalizer**
  Chapter 4: C1 EQ + LLcompressor
  C1 Compressor + EQ

- **Equalizer + Gate**
  Chapter 4: C1 EQ + Gate

- **compressor + noise reducers**
  Chapter 5: C1 HLcompress + DeHiss

- **expander + de-hissing**
  Chapter 6: C1 Noise Reducer

- **compressor + de-hissing**
  Chapter 5: C1 HLcompress + DeHiss

- **compressor + de-esser**
  Chapter 7: C1 Compressor + DeEsser 1

- **compressor + enhancer**
  Chapter 8: C1 HLcompress+bass/treb enhance
  C1 HLcompress+treble enhance
  C1 MLcompress+bass/treb enhance
  C1 MLcompress+treble enhance

- **De-Esser + gate**
  Chapter 3: C1 DeEsser+ Gate

- **Enhancer + gate**
  Chapter 8: C1 Bass/treb enhance+Gate
  C1 Treble enhance + Gate

- **compressor + de-reverberation**
  Chapter 5: C1 Compressor + De-Reverb

In most cases, the compression may either be of conventional high-level type, including limiting, or mid level compression leaving the dynamics of the loudest parts of sounds unaffected, but bringing up quieter details.

On using setups

All the setups in this library should be treated as starting points. You will always have to adjust the setups for use with each individual sound file, to match its loudness, spectral content, dynamics, and specific problems it suffers from, as well as the specific results you desire.

For your convenience this manual describes for each setup which controls are most useful and how they should be adjusted to achieve the desired effect. The actual adjustment uses your skill as a listener to decide the most satisfactory effect, and suggestions and hints are for guidance, not limitations on your creativity or skill.

The guidance for each setup describes not just its purely technical use but a little about what kind of effects and side-effects can be expected so that the user can listen for these and make a judgement as to the best trade-off in any particular case.

In many setups you may in particular have to adjust the output gain to prevent overload clipping on particular signals, due to transient overshoots etc. Instructions on how to prevent overload clipping are given in section 2.5 (Output level and clipping) of this manual.

In all setups, you have the option of using or not using lookahead. This will not affect the basic functionality of any setup but may affect fine details of sounds. For further details on the effects of lookahead see the C1 User's Guide section on Lookahead and section 2.6 (Lookahead) of this manual.

This manual is divided into a number of chapters, several of which are devoted to setups performing a particular kind of task.

These specialized chapters may be considered almost as product manuals in their own right, devoted to describing the use of a specialized "product" obtained when the C1 is loaded with a particular kind of setup.

There is however, considerable overlap between chapters, because many of the setups in effect involve two processing effects working together, so that some degree of cross-reference between chapters is necessary when the "other" effect belongs to another chapter.

In using this setups Library manual, it is assumed that you have the basic skills in using the C1 found in the WaveSystem Plug-in General Controls chapter of the C1 User's Guide. If not, please review this chapter!

Chapter 2 - Basic Use of Setups

Quick tour of the C1

This section is not a substitute for the C1 User's Guide, but simply an overview that may be found useful in using the setups library, since extensive reference is made to the module structure of the C1 in setups descriptions.

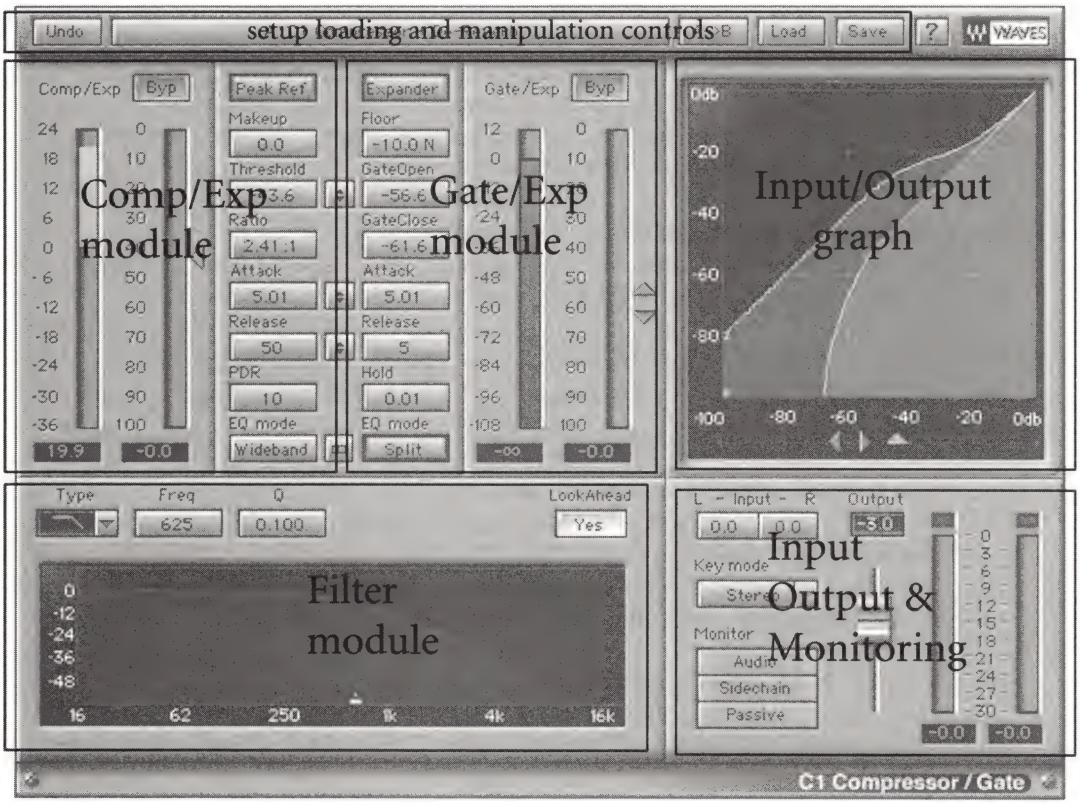

The C1 user interface is divided into the main areas shown in the above figure.

These areas are:

Two dynamics processing modules, and their controls and metering:

- The Comp/Exp ("Compressor/Expander") module
  This acts as a compressor or an expander (which may operate at low, middle or high input levels). It includes two bar meters to monitor respectively gain reduction (in red) or increase (in yellow), and input dynamics control level (in blue). The various buttons and Value Windows control the operating parameters, as described by the captions, and in detail in the main C1 User's Guide.

- The Gate/Exp ("Gate/Expander") module
  This acts as a Gate, or a low-level (downward) expander or low-level compressor. It includes two bar meters to monitor respectively gain reduction (in red) or increase (in yellow), and input dynamics control level (in blue). The various buttons and Value Windows control the operating parameters, as described by the captions, and in detail in the main C1 User's Guide.

- A filter module, including graphical display of filter responses
  This filter module may be used to equalize just the control sidechains of the dynamics processor, or to split the audio into two bands only one of which is dynamically processed, or any combinations of the two. Filter type, frequency and Q buttons and Values Windows are provided for filter adjustment. A graphical display shows the frequency responses of the "active" filter band in red, and the complementary "passive" band of the bandsplit in blue. The filters may also be adjusted by clicking and dragging the cross-marker on the graph.

- An input/output (I/O) graphical display of dynamic signal levels
  This displays the way the output level of the two dynamics modules (yellow for the Comp/Exp module, light blue for the Gate/Exp module) varies with input level, and shows the moment-by-moment variations of level on the graph itself. The triangular grab markers below the graph display, and can be used to adjust, the threshold settings of the two dynamics modules by clicking and dragging.

- An input and output signal level control and monitoring area
  This includes an output level slider, input level controls for left and right channels (which are normally left at 0 dB unless alterations in balance are required), peak-reading output level meters for the 2 stereo channels, and buttons to switch between normal stereo operation and keyed operating modes, and a monitor button to allow auditioning either of the processed audio or of the sidechain or passive bandsplit signal band signals for setting up purposes.

- A number of controls for saving, loading, copying and altering setups

- Controls for previewing and processing sound files

- A title bar indicating name of loaded setup

- An IDR button for optimizing digital resolution at the 16 or 20 bit wordlength.

Detailed instructions on the use of controls are given in the main C1 User's Guide.

Controls which are recommended for use with setups below are indicated within the text in bold typeface.

Loading setups

Place the setup files you wish to use in one (or more) folder of your choice. We shall call this folder the C1 setups folder, but you may choose any other names, and any location on your hard (or other) disc(s) you wish.

Please be sure to keep safety copies of all setup files in the setups Library in another location (e.g. on a floppy disc) so that you have the original setup files in case you accidentally alter the ones in the C1 setups folder(s).

To load the setup you wish to use, click on the Load button, and select the setup file you wish to use, and either click on the Open button or double click on the setup file name. The setup name will appear at the top of the C1 window, and all buttons, Value Windows and graphical displays will be set up to the pre-set settings in that setup.

You may load two setups in the C1 at a time. The SetupA/SetupB button may be clicked to change from one of these setups to the other. If you load a setup in the setup B position, this will not affect the setup loaded into the setup A position, or vice-versa.
This is useful for making quick switched A/B comparisons between two different setups, for example to check whether one sounds better than another.

The name of the last setup loaded or copied into or saved from Setup A or Setup B will be displayed in the title bar at the top of the C1 window.

Making use of Setups

Once you have loaded a setup, you may alter the value or setting of any control on the C1 in the manner described in the main C1 User's Guide.

In most setups, certain controls are important in adjusting the effects and other controls should not normally be used. In the instructions for each setup, the controls that may be adjusted are indicated in bold type.

The simplest way to develop new customized setups for your own use is to modify the values of controls from a single setup in the Setups Library. However, there is another way of compiling new setups from existing ones, i.e. editing by copying and pasting controls settings from one setup to another. Editing allows you to take "bits and pieces" from different setups and incorporate then all into a single new setup. This, for example, you can take some or all of the Comp/Exp module button and Value Window settings from a first setup, some or all of the Gate/Exp module button and Value Window settings from a second setup, and some or all of the filter button and Value Window settings from a third setup.

Editing is most useful when the Comp/Exp and Gate/Exp modules are set up to do completely independent tasks. Then you can take one module from one setup and the other from another setup to compile a new combination setup. However, be warned that if both use the filter module (i.e. if neither are in Wideband EQ mode), this will only work if both use the same filter module setup. If either of the modules is in Wideband

mode, this will not be a problem.

You can do editing by compiling the setup you want in setup A and loading the setups you wish to edit from, one at a time, into setup B.

Hint: To save time, load the first setup you are compiling from straight into setup A. Then you do not have to paste its control settings from setup B to setup A.

Saving and Editing setups

After a setup has been altered, according to the instructions given and your judgement, you may wish to save it for future use on similar sound files. This is done by clicking the Save button. You will enter a dialog box which allows you select a folder in which to locate the setup, or even to create a new folder to save it in, according to the usual Apple conventions. You should type the name you wish to give the setup in the "Save file as:" window, and then click the Save button in the dialog box. After you have done this, the name of the setup in the title bar at the top of the C1 window will change to the new setup name.

Warning: Before saving a setup, ensure that the monitor button is in Audio mode, unless you specifically want the setup to load later in another monitor mode.

The simplest way to develop new customized setups for your own use is to modify the values of controls from a single setup in the Setups Library, or from previously prepared setups you may have. However, there is another way of compiling new setups from existing ones, i.e. editing. Editing allows you to take "bits and pieces" from different setups and incorporate then all into a single new setup. This, for example, you can take some or all of the Comp/Exp module button and Value Window settings from a first setup, some or all of the Gate/Exp module button and Value Window settings from a second setup, and some or all of the filter button and Value Window settings from a third setup.

Editing is most useful when the Comp/Exp and Gate/Exp modules are set up to do completely independent tasks. Then you can take one module from one setup and the other from another setup to compile a new combination setup. However, be warned that if both use the filter module (i.e. if neither are in Wideband EQ mode), this will only work if both use the same filter module setup. If either of the modules is in Wideband mode, this will not be a problem.

You can do editing by compiling the setup you want in setup A and loading the setups you wish to edit from, one at a time, into setup B. To paste the settings of a number of control buttons, value widows and/or grab markers from setup B to setup A, proceed as in the following example.

First, in setup B, select the desired buttons, Value windows and/or grab markers by click and holding down the mouse button at a point outside the desired controls (not in a button or Value Window!), and dragging the mouse to form a rectangle intersecting or including just the desired controls as shown in the following illustration for all the Comp/Exp module controls:

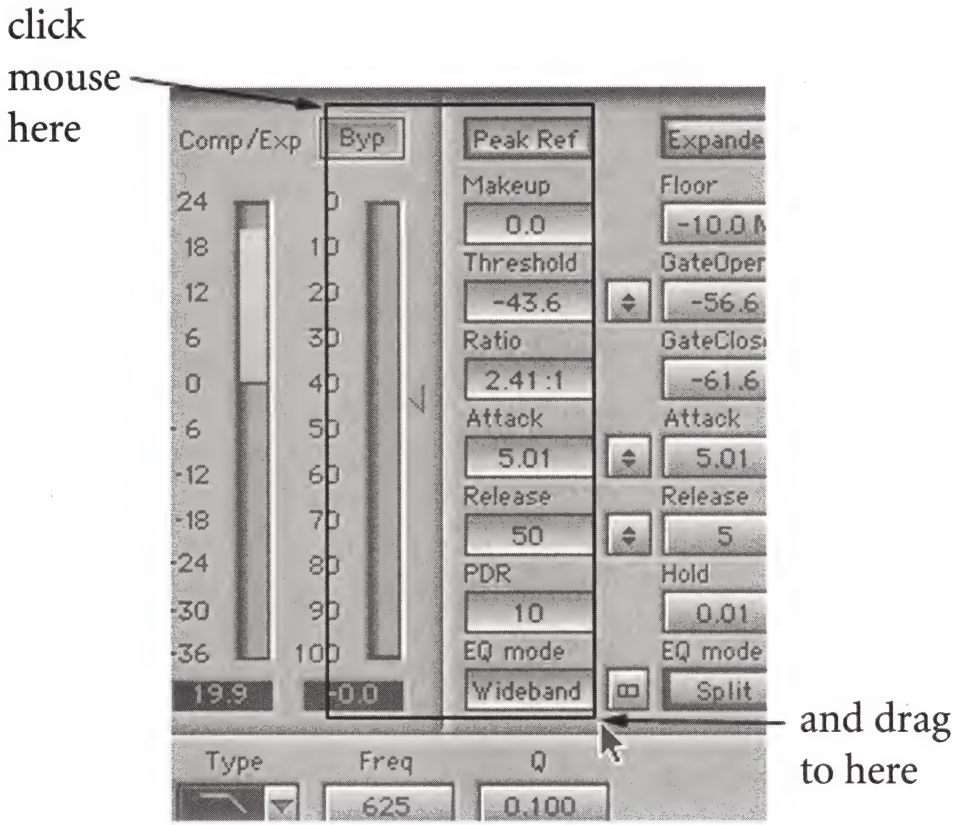

Letting go of the mouse button, you will have selected all controls intersected by the rectangle. Then press the c key on the keyboard (NOT control C !) to copy the control settings. Then switch to setup A by clicking the Setup A button, select the same controls as before by click and drag, and then press the v key on the keyboard (NOT control V!) to paste the control settings.

This completes the pasting operation.

You may switch back to setup B by clicking the setup B button, load a new setup from the setups Library into setup B, and paste control setting from another part of this setup into setup A by repeating the above procedure. And so on until you have compiled your new setup in setup A. Then you can save it if desired, or just use it for processing.

The same method can be used to compile a setup into setup B from setups loaded into setup A.

Hint: To save time, load the first setup you are compiling from straight into setup A. Then you do not have to paste its control settings from setup B to setup A.

Output level and clipping

You may sometimes find on particular sounds that use of setups Library setups will cause clipping or overload distortion. Such clipping can be caused by excessive gains or from transient overshoots.

The internal processing in the C1 is designed so that clipping distortion cannot occur in the internal signal processing algorithm, so that the only risk of clipping comes from excessive output gain. Thus all you need do to prevent clipping is to reduce output gain appropriately.

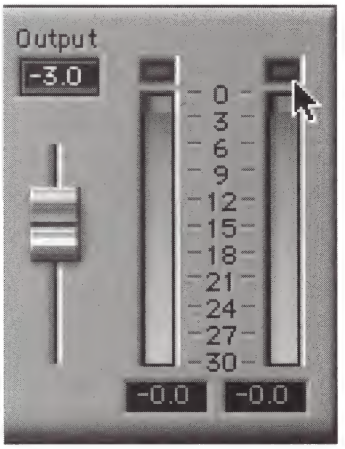

Output meter showing ten clippings and with mouse positioned to reset the overload indication by clicking.

The windows above the output level meters show the number of overload clippings occurring during preview or during processing a file in red. They may be reset before a preview or processing by clicking on them. If, after processing, clipping is shown, you may wish to undo the processing and to re-do it with a lower output gain to prevent clipping, although one or two or very few clippings only may prove not to be audible.

The readings at the bottom of the output level meters shows the maximum peak level achieved during preview or processing. It may be reset before preview or processing by clicking on the meter itself. These readings are useful to allow increase in output level if the sound is undermodulated. If for example, the peak level shown in previewing or processing a file is -4.7 dBFS (relative to digital full scale) in the highest of the two channels, then you may increase the output level by say 4.6 dB without the re-processing of the file (after undoing previous processing) causing overload.

However, be aware that the peak level occurring in processing a whole sound file may not be indicated during preview of just a limited segment of that file. You can check peak levels either by

(i) processing the file, and then undoing the processing after taking meter readings, or

(ii) clicking the playback button so that it is illuminated, and then playing the entire sound file under Sound

Designer II™, or at least all parts of the sound file which are loud enough to be expected to cause clipping problems.

Hint: Especially when using compression or limiting, you may find that transient peaks are reduced in level, often by 2 or 3 dB, if you process using Lookahead in Yes mode, rather than process without Lookahead. However, this can also subtly affect the sound, so use your judgement here as to whether this is desirable in your case.

Lookahead

The use of the lookahead button will not seriously affect any of the setups in the setups library and may be set to "Yes" or "No" as desired by the user.

To some users, the function of the lookahead button may seem mysterious, as it at first sight seems to do little. Internally, it delays the audio signals just enough to match inevitable delays required to derive the gain control signals used to alter signal gains in compression, expansion and gating. Most analog dynamic processors do not have lookahead. The C1 provides the option of simulating analog processing by switching lookahead out (i.e. the lookahead button displays No). With lookahead in (i.e. the lookahead button displays Yes), generally, the shape of compressed transients is better, with less overshoot (see the illustrations below), and there is less premature gating of initial transients such as sibilants or drumstick sounds.

4 kHz test signal with sudden 20 dB step

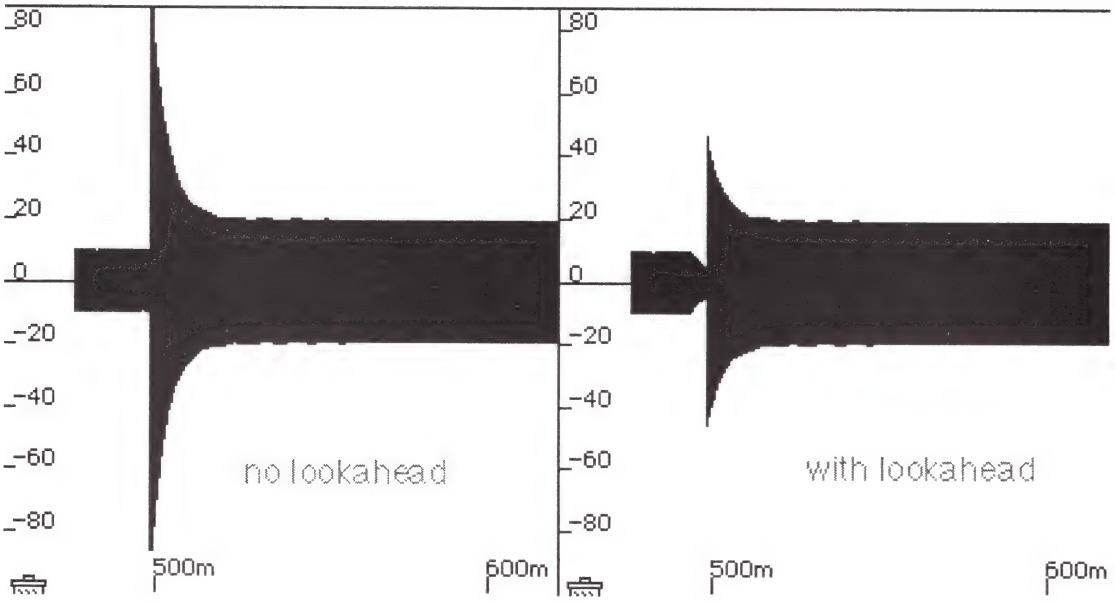

Effect of compression on 4 kHz test signal without and with lookahead for limiting at threshold -30 dBFS with 5 msec attack time. Note the reduced overshoot of the transient. The effect is less extreme on most real world audio waveforms!

Thus in most situations, the processing is better behaved with lookahead, and setups in the setups library are loaded with lookahead.

However, the actual sound without lookahead is different - for example, transients often sound brighter if compressed without lookahead, so this may be preferred on the basis of sound. The initial gating of transients like short sibilant sounds may also sometimes be wanted to modify the sound - and again no lookahead would then be preferred.

Also, because lookahead starts modifying the sound slightly before the start of a transient, it can sometimes cause a subtle but noticeable "pre-echo" effect sounding like a slightly disturbing room echo, especially on high-quality speech. If this effect is audible, you may again prefer not to use lookahead.

As seen in the above waveform illustrations taken from Sound Designer II™, the time delays used to implement lookahead have no effect on the timing of processed files, which remain exactly synchronized to unprocessed files.

Chapter 3 - Basic Dynamic Processors

Classic processors

This chapter describes a few setups that are what may be termed "classic" dynamic processors, the kind of basic workhorses used for most dynamic processing in the audio industry: the Compressor/Gate, Compressor/Expander and DeEsser/Gate.

Classic Compressor and Gate

This is a classic (high-level) soft-knee compressor/limiter with a classic gate.

Setup name: C1 Classic Compressor/Gate

The Comp/Exp module is here set up as a wideband high-level compressor, and the Threshold, Ratio, Attack and Release controls are conventional in use, provided that the ratio is kept within the range 1 to 50. Above a ratio of 20, the compressor acts as a soft-knee limiter. The Makeup gain allows extra gaiun to compensate for any loss of level due to the compressors reduction of high-level gain.
The PDR (program-dependent release) allow the compressor to respond more rapidly to transients of limited duration, giving a faster release for short transients. The setting here is the duration in msec of transients for which the release time is more rapid. For morte sustained sounds of longer duration, the release time is that set by the Release control. PDR minimizes prolonged gain reductions caused by short transients.

For A/B comparisons of the processed with the original signal, the gain of the bypassed signal may be adjusted by using the bypass gain value window (below the bypass button) so that there is little audible gain change when switching between bypass in and out.

The Gate/Exp module is set to a classic fully-functioned gate. As loaded, the GateOpen and GateClose levels are pre-set 4 dB apart to minimize gate "chatter". The hold control ensures a minimum gate on time which again helps prevent chatter. Attack and Release may be adjusted in the usual way. If it is desired to retain some of the background atmosphere during quiet passages, the floor control may be used to adjust the remaining signal level when the gate is "off".

The Compressor may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The Gate may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

Classic Compressor and Expander

This is a classic (high-level) soft-knee compressor/limiter with a classic soft-knee 2:1 downward expander.

Setup name: C1 Classic Compressor/Expander

The compressor is identical to that in the previous setup, and used in the same way.

The Gate/Exp module is now in low-level (downward) expander mode. GateOpen controls the expander threshold. Attack and Release may be adjusted in the usual way. If it is desired to retain some of the background atmosphere during quiet passages, the floor control may be used to adjust the remaining signal level. The expander is a soft-knee 2:1 device, but the effective expansion ratio below threshold may be increased by using the "negative" polarity settings of the floor control (indicated by "N" after the dB indication in the floor window), with increased effective ratio as the floor setting moves between -100 N and -10 N.

The Compressor may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The expander may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

Classic De-Esser and Gate

This is a classic de-esser with a classic gate.

Setup name: C1 Classic DeEsser + Gate

The Comp/Exp module here is used to implement a de-esser to reduce the level of high-level sibilants. This is implemented as a limiter responding to the level of an equalised sidechain signal in the "Ess" frequency band. The attack and release times are intentionally short to minimise the effect on the signal once an "Ess" sound is over. The degree of de-essing ixs adjusted by moving the threshold up or down. The only other controls used for de-essing may sometimes be the frequency to tune the precise band of "Ess" frequencies.

The gate is identical to that in the C1 Classic Compressor/Gate setup above.

The De-Esser may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The Gate may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

About Compression

Compression is the process of reducing the dynamic range of a sound, by altering the gain. So, for a given number of dB change of level at the input, a smaller number of dB change of level takes place at the output. The ratio

$$\frac{\text{Input dB change}}{\text{Output dB change}}$$

is termed the compression ratio, so that if every 2 dB input change causes a 1 dB output change, the compression ratio is 2.

Conventional "hard knee" compressors normally have a constant compression ratio when the input signal rises above a user-preset "threshold level, but the C1 has a more sophisticated compression characteristic matched better to the way the ears hear sound, rather than to an abstract mathematical law. This is not only a "soft knee" law - designed to be neither too hard nor too soft, but involves other features allowing uses in different kinds of compression modes.

The C1 allows three distinct kinds of compression of sounds -

- conventional high-level compression, for which the highest sound levels above a threshold are compressed and reduced in level.
- mid-level compression, in which sounds in the middle of the dynamic range are compressed, but where both very low level and very high level sounds are not compressed.
- low-level compression, for which the very quietest parts of a sound are raised in level to make them louder or more audible, but the louder parts of a sound are unaffected.

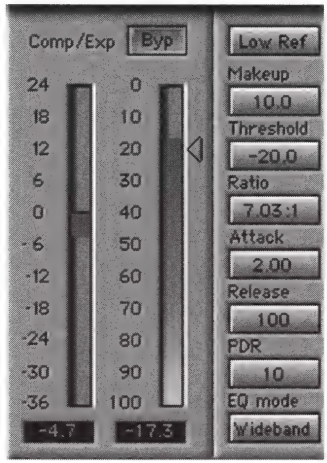

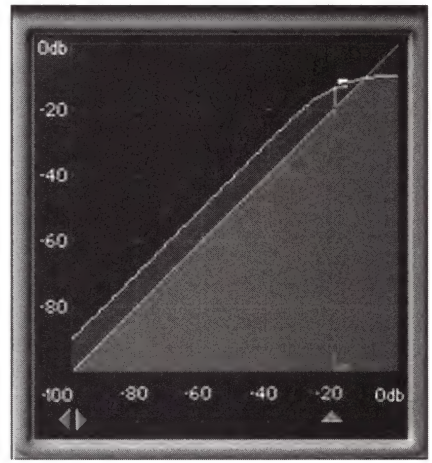

Typical high-level compressor setup.

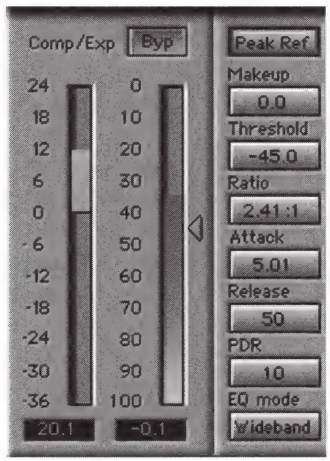

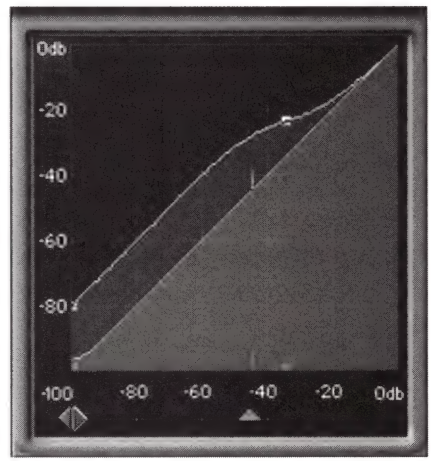

Typical mid-level compressor setup.

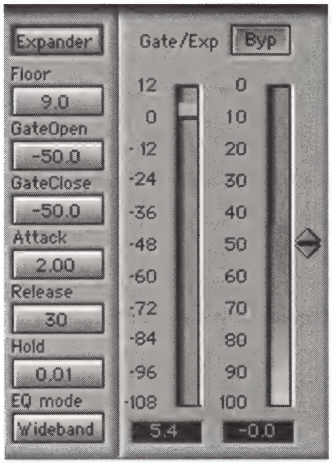

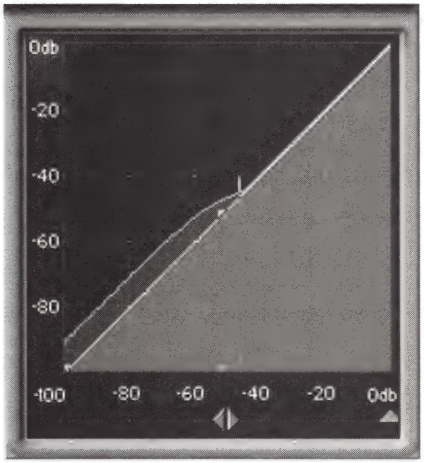

Typical low-level compressor setup.

-These three kinds of compression all have important uses, and this flexibility of the C1 makes it a much more effective tool for obtaining precisely the kind of compression needed in any given application. Using conventional high level compression in an attempt to solve all problems often results in an over-compressed sound which has "all the life squeezed out of it", whereas choosing the appropriate form of compression can result in a much more pleasant and natural sound.

Setups for all kinds of compression in the C1 are available. In each case, the other dynamic module can be used for other tasks, often to solve problems that arise when compression is applied. Compressed sounds often have excessively audible low-level noise, or other low-level sounds such as room reverberation or echoes can be increased to an annoying level. The other processing module can often be used to counteract these problems.

Setups in the Library using the three types of compression include:

High-level compression:
C1 Classic Compressor/Gate
C1 Classic Compressor/Expander
C1 Classic DeEsser + Gate
C1 Compressor + EQ
C1 HLcompress + DeHiss
C1 HLcompress+bass/treb enhance
C1 HLcompress+treble enhance
Mid-Level compression:
C1 Multimedia Speech 1
C1 Multimedia Speech 2
C1 Speech Compress/Expand 1
C1 Compressor + De-Reverb
C1 Compressor + DeEsser 1
C1 MLcompress+bass/treb enhance
C1 MLcompress+treble enhance

Low-Level compression:
C1 EQ + LLcompressor
C1 De-Ess + LLcompress

In general, wherever a wideband high-level compressor is used in the Comp/Exp module, it may be replaced by a mid-level compressor, and vice-versa, by pasting Comp/Exp control settings from one setup to another.

Which type of compression?

The choice of type of compression depends on what you wish to achieve, and also the nature of the sound file.

If the aim is simply to get the loudest sound possible, then high-level compression or limiting should be used. But be aware that this can also cause the compressed sound to lack a sense of dynamics and seem oppressive. Additionally, especially over cheap reproducing systems with small loudspeakers (e.g. portable radios, portable headphone stereos and small computer loudspeakers), the high average levels can drive the amplifiers or loudspeakers into distortion, giving a fatiguing and unclear sound for the listener.
If however, the aim is not loudness for its own sake, but to ensure that the quietest passages are clearly audible even under difficult or noisy listening conditions (e.g. in an automobile) or via systems with a high floor noise or distortion level (e.g. 8 bit multimedia files), then what you may need is low-level or mid-level compression. This raises the quieter passages above the low level at which they are not adequately audible. This will make those passages clearly audible and less prone to low-level noise or distortion without losing the sense of high-level dynamics, and without pushing all louder sounds to near peak level.

Mid-level compression combines these virtues with making the wanted sound seem louder, by raising the level of the middle sound levels at which signals such as speech spend most of the time.

Low-level compression is preferred where you do not necessarily wish to make the louder sounds in a sound file louder still, but where you need to make the quietest passages or details of sound more audible. Classical music, or drama or documentary where background ambiences need to be made more audible, particularly can benefit from low-level compression.

Chapter 4 - Simple Setups

Multimedia Speech setups

These setups are wideband compressors + gates or downward expanders suitable for use when preparing multimedia speech files, especially those that will later be reduced to 8 bit resolution. This setup may also be used with other multimedia sounds and with speech in other applications where loudness must be combined with quietness between words.

With 8 bit multimedia files, the 8 bit quantization introduces several quality problems. If the sound is above the quantization "floor" level, the quantization sounds like a steady hiss, and it is desirable to minimize its subjective level by maximizing speech loudness. However, if this is done, any background noise is brought up in level, and is liable to give cause the quantization noise to sound very unpleasant between words, with a distinctly "grainy" and intermittent quality. The solution here is to combine mid-level compression to bring up speech levels to be above the quantization noise level during syllables, with gating or effective downward expansion of low-level input signals to prevent them causing the granular sound of bad quantization noise. To minimize loss of quality, this processing should be done on 16 bit files on the C1 possibly at a 44.1 or 48 kHz sampling rate, later using a specialist tool such as the Waves L1 Ultramaximizer to perform the final stage of conversion after sampling rate conversion to maximize file levels and convert to 8 bits.

Soft-knee mid level compression of the kind unique to the C1 is ideal for multimedia speech, since it compresses the middle speech levels at which speech spends most of its time, without "squashing" the dynamics of speech peaks or of low-level sounds. In this way, a more natural loud sound is obtained than by using conventional high-level compression. The problem that quiet sounds are brought up a great deal in level is solved by the gating or downward expansion.

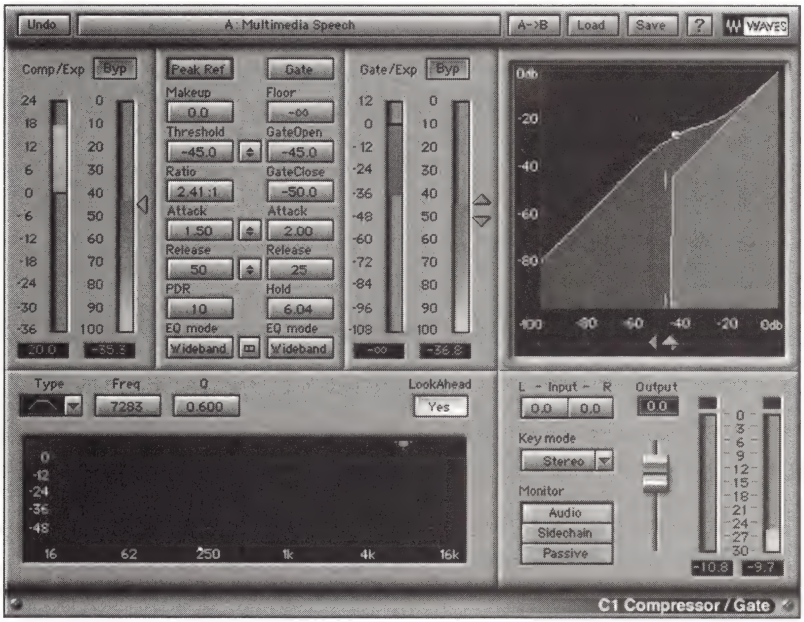

Setup name: C1 Multimedia Speech 1

The Comp/Exp module is set up as a mid-level compressor to squeeze the top 60 dB of the input dynamic range into the top 40 dB of the output dynamic range, as can be seen from the yellow input/output curve on the graph, This causes a large gain increase at low levels as can be seen from the yellow bar in the Comp/Exp gain reduction meter. The threshold and ratio (between perhaps 1.5 and 3) may be adjusted to vary this compression curve, using the "PeakRef mode" to keep the top of the curve in the right place.

The high gain of low level sounds thus caused (20 dB for the setup illustrated above) may be counteracted by using the Gate/Exp module as a gate to switch off sounds below a certain level. The gate is a conventional one with GateOpen and GateClose thresholds, Attack, Release and Hold adjustments. The values shown are typical of those useful in gating speech, although each case will require different adjustment. For multimedia 8 bit applications, a floor of minus infinity as shown should be used, and the GateOpen thresholds should not be much lower than shown to avoid audible inter-syllable quantization granularity in the final 8 bit sound file.

For gating speech with attack times such as 2 msec or even 5 msec, initial speech sibilants tend to get diminished or cut off unless you use lookahead in the "Yes" position.

Setup name: C1 Multimedia Speech 2

This setup is not illustrated because it is identical to the above except that the Gate button is switched to "Expander" mode (this may be done by clicking on it). This turns the Gate into a downward expander, whose adjustments are similar except that GateClose and Hold are inoperative. The downward expander, used at an extreme "negative polarity" floor setting of -10 N, gives a useful degree of reduction of background noises, and may often be preferable to a gate on account of its gentler "soft knee" reduction of gain.

Equalizer with Wideband dynamics processor

Among the simplest setups to understand are those that combine a single conventional wideband dynamics processor with a simple equalizer, allowing both dynamics processing and EQ to be carried out in a single step. The C1 allows one of its modules to be used for dynamics processing, while the other one is used in a mode where no dynamics processing occurs, but where the use of "split band" plus a suitable "makeup gain" in that band gives equalization.

Setup name: C1 EQ + Gate

The Gate/Comp module in this setup is a normal "wideband" gate, and it may be adjusted in the usual way for gates. It is here operated in "lookahead" mode to minimize loss of starting transients.

The Comp/Exp section is set to ratio 1 so that it does absolutely no processing, but it is in split band mode, which allows the gain of the active frequency band (shown in red on the following graph) to be raised or lowered.

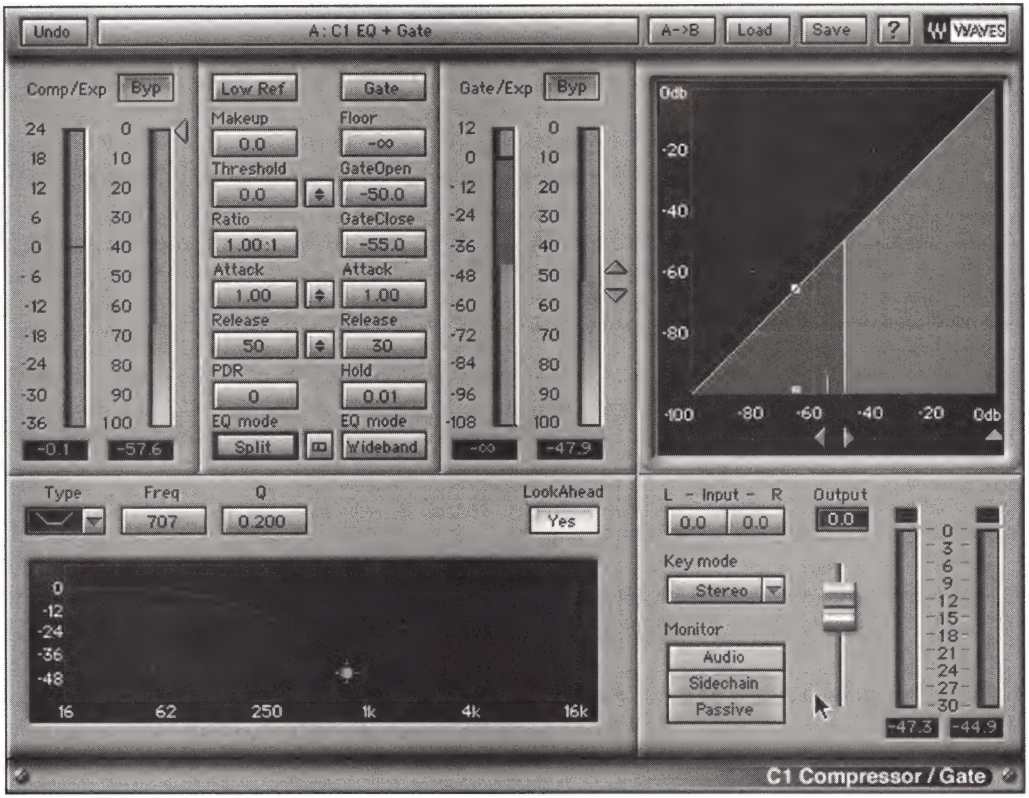

Adjustment of the EQ in no way affects the operations of the gate. The two processes are completely independent.

The equalizer gain in the active band is adjusted by click-and-dragging the Makeup gain Value Window (highlighted in the previous graphic). The filter type, frequency and Q controls may be used to adjust the shape of equalization used. As loaded, the bass below 125 Hz and treble above 4 kHz is boosted (for Makeup gain greater than 0 dB) or cut (for Makeup gain less than 0 dB). The is achieved by using filter type bandreject, which attenuates middle frequencies but lets through low and high.

Other filter type settings obtained by clicking on the type button are:

lowpass, which may be used to boost or cut lower frequencies,
highpass, which may be used to boost or cut higher frequencies, or
bandpass, which may be used to boost or cut a band of frequencies centered at the preset frequency.

The EQ may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The Gate may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

Setup name: C1 EQ + LLexpander

This setup combines an equalizer with a conventional wideband low-level (or downward) expander.
The equalizer is used as in the C1 EQ + Gate setup, and the low-level expander is implemented by the Gate/Exp module, whose controls are used in the usual way. This setup may be used for example when the equalization brings up background noises, where the expander may help to reduce them again without having the sudden switching effect that can be caused by a gate.

The EQ may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The LLexpander may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

Setup name: C1 EQ + LLcompressor

This setup combines an equalizer (implemented and adjusted as in the previous two examples by the Comp/Exp module makeup gain plus the filter module) and a low-level compressor implemented by the Gate/Exp module.

The Low level compressor brings up the level of low-level sounds below a threshold set by the GateOpen control. The amount of boost, between 0 and +12 dB, is set by adjusting the Floor control. This effect is useful to make inaudible low-level details in a sound more audible, whether these are background ambient noises needed for "atmosphere", quiet from-audience speech at a low level in a conference recording, or quiet noises associated with a louder sound whose increase in level will make the sound seem "hyper-real".

The equalization, adjusted as before, does not affect the low-level compression dynamic processing, but only the tonal quality of the result.

The EQ may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The LLcompressor may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

Setup name: C1 Compressor + EQ

Unlike the previous setups, this has the Comp/Exp module used for wideband dynamic processing and the Gate/Exp module set up in split mode with the filter as an equalizer.

The equalizer is adjusted exactly as before, except that now the floor control is used to adjust the EQ gain rather than the Makeup control. So the floor control is raised to increase gain in the active band and lowered to decrease it.
The Comp/Exp module is adjusted as a high level compressor in the conventional way with threshold, makeup gain, ratio, attack, release and PDR (program dependent release) controls adjusted in the usual way, and as described in the C1 User's Guide.

The Compressor may be switched in or out by clicking on the Comp/Exp bypass button (over the Comp/Exp bar meters).

The EQ may be switched in or out by clicking on the Gate/Exp bypass button (over the Gate/Exp bar meters).

The Comp/Exp section may also be set up as a mid-level compressor as used in the multimedia speech setups of section 4.1 above.

compressor setups

Here we describe a few more setups combining compression with other functions. For a basic review of compression see sections 3.5 and 3.6 above.

Speech compressor/expander 1

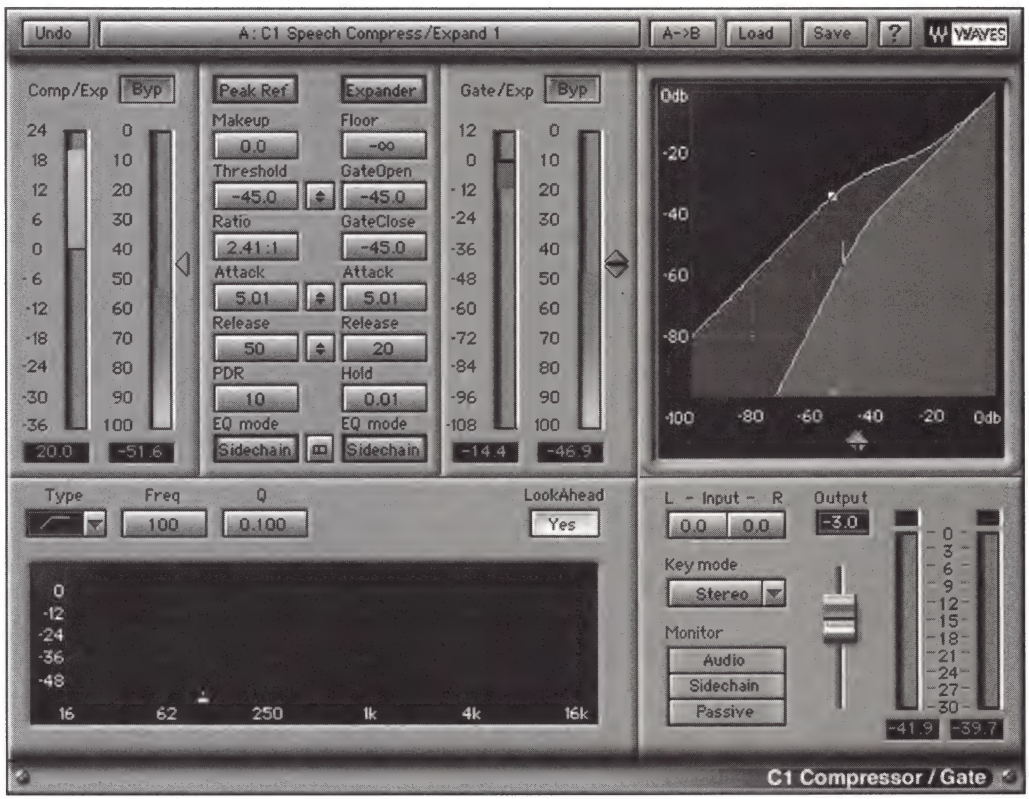

Setup name: C1 Speech Compress/Expand 1

This setup allows a high degree of speech compression and increased level with minimum effect on the naturalness of speech sound. The more natural compressed quality has the downside that the objective unweighted dynamic range can vary more widely than with a wideband compressor.

One important feature is the use of a sidechain EQ that rolls off gently in the bass - note the low Q. (The actual sidechain EQ in dB is 1.5x the number of dB indicated). This reduction of bass gives a much better and more natural compressed sound on speech, and a much less "constricted" sound. Typically the best effect occurs for frequencies from 60 to 180Hz.

The Comp/Exp module is used as a mid-level compressor, and the ratio and threshold controls can be adjusted to taste over a range in which the ratio may be varied between 1.5 and perhaps 5 and the threshold between - 30 and say -70 dB.

Although the sidechain EQ gives a more natural sound, the high degree of compression brings up background noise and reverb. This may be countered by the low-level expander of the Gate/Exp module, set to come into action at the bottom of the compression curve, and also using the sidechain EQ.

This has the effect of reducing noise and also any room reverberation present.

GateOpen is the threshold for downward expansion of noise and reverb, and a setting as high as possible that has little or no audible effect on speech level is recommended, providing that this does not cause too much audible modulation of noise.

One may drag all thresholds (for Comp/Exp and Gate/Exp module) together by click-and-dragging the button between the threshold and GateOpen Value Windows when adjusting threshold to particular speech and background noise characteristics.

The compressor may be bypassed by clicking on the Comp/Exp button, and the expander may be bypassed by clicking on the Gate/Exp button.

Mid-level compression plus speech De-reverberation

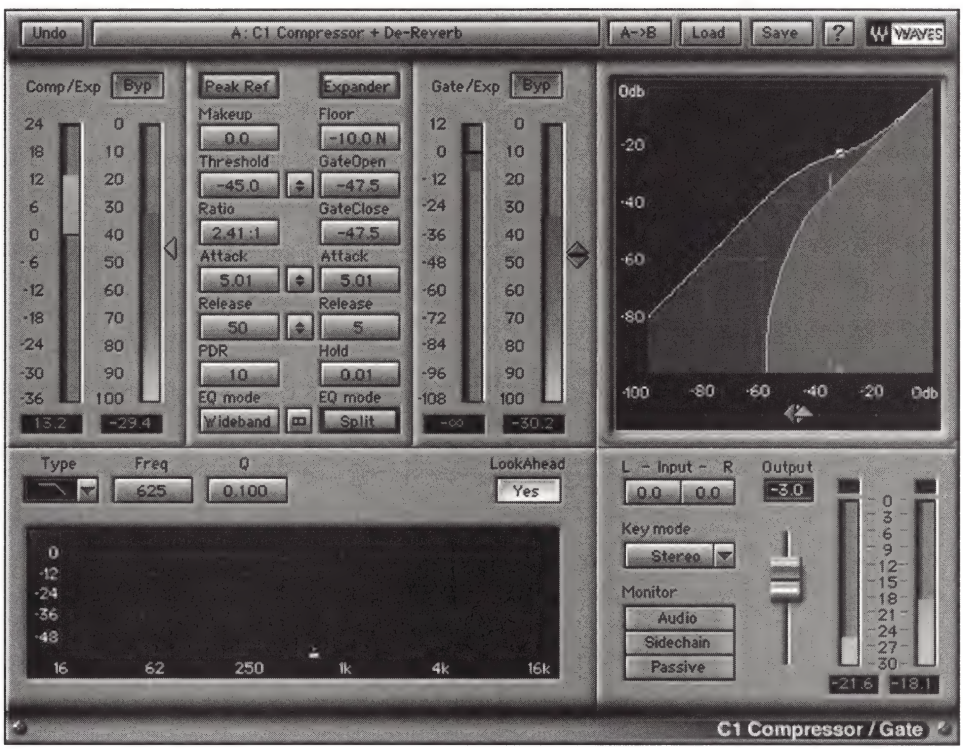

Setup name: C1 Compressor + De-Reverb

This setup uses the Comp/Exp module as a mid-level compression unit and squeezes the top 60 dB of dynamic range into 40 dB without compressing the lowest or highest level sounds. Such mid-level compression is useful for making sounds louder and for fitting sounds into the limited dynamic range of say 8 bit multimedia files, without causing the sound to drive peak levels so hard that the sound becomes unpleasant on cheap systems.

Such mid level compression, however, brings up low level parts of the sound by around 20 dB and this can cause unpleasant side effects. The Gate/Exp module can be used in various ways to reduce these side effects.

One side-effect on speech is that room reverberation, which can often already be unpleasant with recordings made outside good studio settings, is further increased to unacceptable levels. The Gate/Exp module here is set up to remove bass signal components at lows levels. This removes background rumble noises and bass components of reverberation tails. Especially on male speech, this has the subjective effect of reducing or removing room reverberation, while having very little effect on speech tonal quality providing the GateOpen threshold is set carefully by ear.

To switch mid level compression in or out, click on the bypass button over the Comp/Exp module bar meters.

To switch speech de-reverberation in or out, click on the bypass button over the Gate/Exp module bar meters.

Controls for mid-level compression:
Comp/Exp module used for mid-level compression.

- Makeup. as needed
- Threshold. -70 to -20.
- Ratio 1 to 3
- Attack, Release, PDR: As needed

Controls for bass de-reverberation:
Gate/Exp module set to bass de-reverberation

- GateOpen threshold. -40 to -15 set by ear to leave signal bass balance unaltered while diminishing reverb bass.
- Attack, Release: As needed - usually attack not too short (to avoid distortion of the bass frequencies) and release not too long (to allow quick response to low level in the bass)

Other controls

- Lookahead
- Frequency: In range 200 to 600 Hz for best sound
- Output Level (has same function in this setup as makeup)

High-level Compressor + De-Hisser.

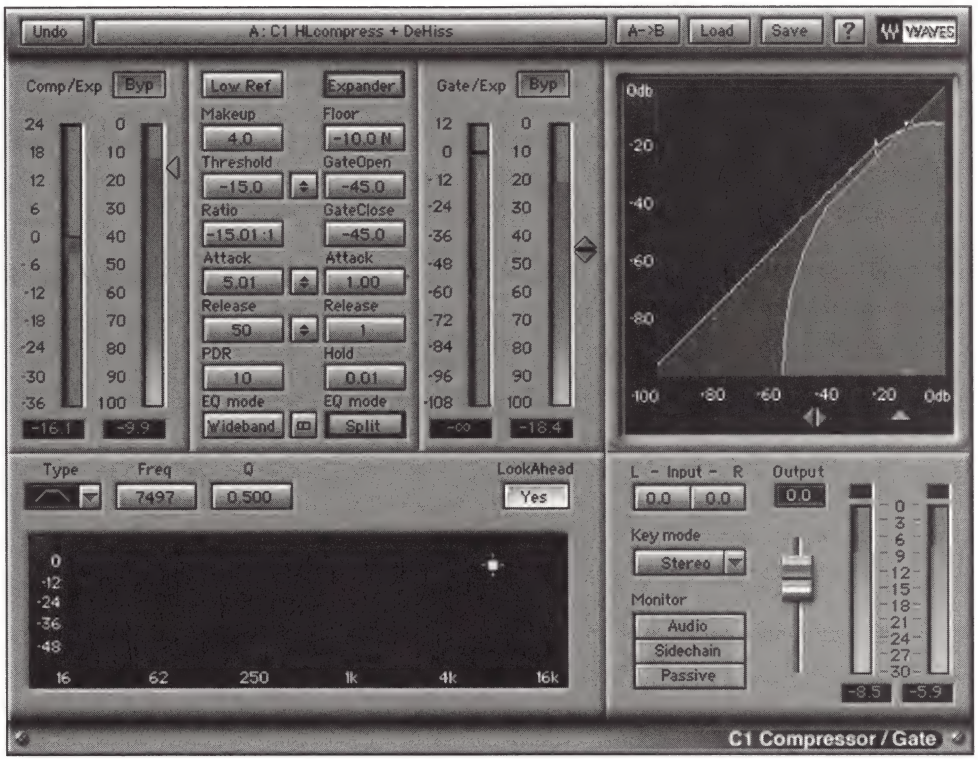

Setup name: C1 HLcompress + DeHiss

This is two effects which may be used independently of each other (by bypassing the other) or together. Each may be set-up and optimised independently - adjustments of one do not affect the other.

The Comp/Exp module is set up as a wideband High Level (HL) compressor with Makeup gain. Note that the Comp/Exp ratio is set beyond the "infinite" setting of full limiting to a negative "over-limiting" ratio setting to bring down peaks a little further. This is very effective providing any cancellation nulls lie beyond 0 dB.

The use of compression with gain has, however, the side effect of bringing up any background hiss in a recording. The Gate/Exp module is set up as a de-hisser to minimise the effect of his with least effect on the sound. It is a low-level expander whose action is confined to a carefully-tuned frequency band centered around 7.5 kHz. (See chapter 6 for more information on noise reduction.)

The filters chosen here for de-hissing are generally suitable for hiss with a white, pink or blue spectrum, requiring only perhaps a couple of kHz up or down tuning of frequency - beware however that one is not tuning out hiss emphasis due to a monitor speaker coloration! The tonal effect of the maximum available hiss reduction can be heard by setting the Monitor button to passive mode. The Q setting turns out to be fairly critical - lower will filter out more hiss, but modulation noise - variations of the hiss up and down with the wanted signal, will become

more audible. A Q of 0.5 is about the best compromise, although a Q of up to 0.6 can be used. Generally, the optimum filter settings do not have a lot of leeway and end up close to those in the setup.

The GateOpen threshold setting is the critical one for de-hissing; it should be adjusted up and down until the highest setting that has no significant unwanted effect on tonal quality of the signal is found. This is the optimum de-hiss setting. The setting is a trade-off between loss of treble on the wanted signal and hiss level. On most speech and popular music without very wide dynamics, there is usually a good setting that has very little tonal effect but a subjective hiss reduction of 6 dB or more. The above GateOpen setting was for a white hiss at a level of only -45 dB relative to digital peak - very hissy indeed, but had little tonal effect.

Generally speaking, de-hissing is more difficult on material with a very wide dynamic range, since the tonal effects in quiet passages will be more audible.

Useful Controls:

Comp/Exp module set to high-level compression/limiter.

- Makeup. as needed
- Threshold. -30 to 0.
- Ratio -50 to -5 or 1 to 50 as needed
- Attack, Release, PDR: As needed

Gate/Exp module set to de-hissing

- GateOpen threshold. -25 to -80 set by ear to leave signal treble balance unaltered while diminishing hiss.

Other Useful controls

- Frequency: In range 5000 to 10000 Hz for lowest perceived hiss
- Q: In range 0.45 to 0.6 for best trade off of modulation noise versus hiss reduction
- Output Level (has same function in this setup as makeup)

Chapter 6 - C1 Noise Reduction

About Dynamic Noise Reduction

The C1 becomes a very powerful and effective single-ended noise reduction system when loaded with Noise Reduction, De-Hiss or De-Rumble setups. It allows reduction or removal of hiss, rumble or other background noises with the minimum of side-effects.

Noise can usually be reduced or removed from a sound file by heavy filtering of the most audible noise frequencies, but this usually also severely degrades the tonal quality of the wanted signal. Another conventional method of reducing noise is to gate or downward-expand low-level sounds, but this also often has the side-effect of altering or removing wanted low-level sounds that are not in the noise frequency range - and one can usually hear the noise coming and going as signal level alters..

In noise reduction setups, the C1 acts as an intelligent filter that only filters noise when it has to, leaving the wanted signal unfiltered when the wanted signal is loud enough to mask the noise in the frequency band being noise-filtered. It combines the virtues of a noise filter and of a downward expander or gate, while minimising the weaknesses of both.

The principle of dynamic filtering of noise, filtering the signal only when it is below a threshold level in the frequency band of greatest audibility, is not a new one. But in the C1 implementation, the frequency band shape, the nature of the dynamic and sidechain filtering, the attack and decay time constants, the expansion law and the nature of the filter as it dynamically varies have all been optimized for optimum subjective results.

Understanding that the noise reduction is a filter that comes in or fades out as the signal energy varies in the noise band is important. Too much filtering, and the wanted signal is affected. Too little filtering, and a lot of noise remains. The adjustment of the noise reduction setups is to obtain the optimum compromise between too much filtering and too much noise.

But it is also important to be aware of a potential side effect - noise modulation. As the filter comes in or goes out, one may hear also the effect of noise modulation, i.e., the noise level going up and down with the wanted signal. This can be a very distracting effect, and sometimes it is less distracting to leave some noise in than to suffer from excessive noise modulation.

All available forms of noise reduction suffer from audible side effects, and require the user to listen carefully, and to adjust them to minimize these. We believe that the C1 offers the best trade-offs among available dynamic filter noise reducers, having fewer side effects for a given degree of noise reduction.

The use of the C1 noise reduction is now described with reference to the setups provided in the C1 setups library. You should load the setups described and try them out on appropriately noisy soundfiles, adjusting them as described below.

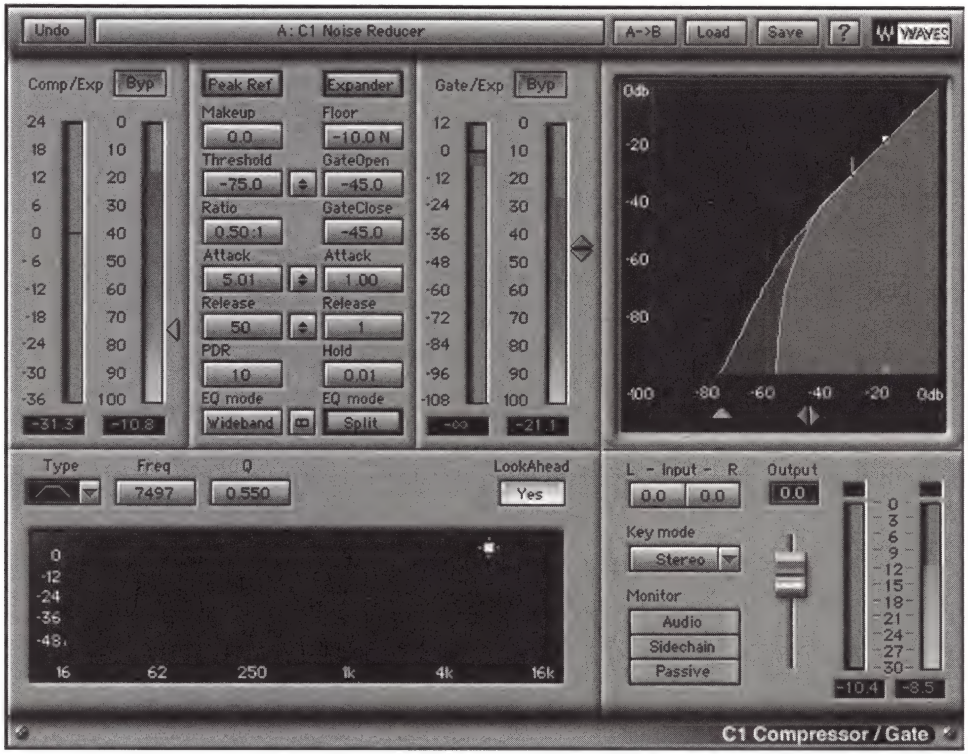

Setup name : C1 Noise Reducer

This is intended for Hiss reduction. It provides two stages of processing. The dynamic noise filtering is provided by the Gate/Exp module, used in Bandsplit mode with the filtering shown on the frequency response graph. The band shown in red is the band whose level is reduced in quiet passages to reduce hiss.

An optional wideband low-level expander is provided by the "Comp/Exp" module. This can, optionally, be used to gently pull down further the level of any very quiet passages in which the reduced noise may still be audible. It can be switched in or out at the user's option by clicking on the Comp/Exp button to bypass the Comp/Exp module. The noise reduction works very effectively on its own, and this low-level expander is provided simply to provide a further reduction in quiet passages if desired.

The adjustments to the Noise Reducer (Gate/Exp module) are essentially just three controls:

(i) GateOpen threshold in the Gate/Exp module.

Moving this upward will increase the amount of filtering both of the signal and noise, moving it down will decrease the amount of filtering. Normally, one would use the highest setting that does not audibly degrade the wanted signal or cause objectionable modulation noise. The precise setting is a judgement depending on the skill and aims of the user.

Hint. It is often best to adjust the GateOpen level by click-and-dragging the button (highlighted in the previous picture) between the threshold and GateOpen Value Windows. This has the effect of moving both GateOpen and Comp/Exp thresholds up and down together. The reason for this is that if you decide to use the low-level expansion option for further reduction of noise in quiet passages, then the level at which this is effective remains matched to the noise reduction GateOpen threshold.

(ii) Filter frequency.

This adjusts the center frequency of the frequency band subjected to filtering. Normally, this would lie somewhere between 5 and 10 kHz depending on the tonal balance of the hiss, with 7.5 kHz being a typical frequency for "white" hiss noise.

(iii) Filter Q.

This adjusts the width of the filtered frequency band. The lower the Q, the wider the band and greater the amount of perceived noise reduction, but lower Q's also increase the risk of audible modulation noise effects. In practice a Q of between 0.5 and 0.6 is usually best for this reason.

Hint: When adjusting the filtering, you may like to switch the Monitor button to Passive mode. This monitors the "passive" band (shown on the frequency response graph in blue), which corresponds to the maximally filtered signal. This allows you to hear the effect of the filtering on its own without the dynamic processing, and is useful for tuning for minimum hiss audibility.

Switching between "Passive" and "Audio" monitor mode, by clicking on the Monitor button also provides a direct comparison between conventional filtering and the noise reduced audio. In many cases, it will be found that the noise reducer mode, once adjusted, almost magically gives the same noise reduction as the Passive mode, but with little of the treble loss given by simple filtering.

Other controls that can be used.

Normally the above controls are all you need ever use for noise reduction adjustment, but a finer control is possible if one has extra patience. The Floor control can be moved from its preset -10 N setting to less extreme "negative polarity" settings between -10 N and -100 N or to floor settings below say -6 dB. This gives a reduced noise reduction effect, but also reduces modulation noise. A combination of floor and GateOpen threshold may give better results than GateOpen threshold alone, especially where only a partial degree of noise reduction is required and where the noise background fluctuates in level. However, juggling the effect of these two controls is more difficult than just using GateOpen threshold.

When the wideband expander is used, its effect may be adjusted by operating its threshold independently of that of GateOpen, although the preset 30 dB difference in values is generally roughly right. The Comp/Exp module's attack, release and PDR values may be adjusted in the usual way for an expander for optimum sound.

In general, you should usually not alter the attack and release time constants from their preset values in the Gate/Exp module - the values used are about optimum for hiss filtering applications, and also quite well matched to high frequency "crackle" and 78 rpm surface noise.

The use of high-pass filter mode for hiss filtering is generally less effective than the bandpass filtering shown above, mainly because it gives more audible hiss noise for a given degree of tonal alteration, being less well matched to the ear's sensitivity to hiss.. However, if it is used, generally use Q = 0.6 to minimize modulation noise effects.

This kind of setup can also be effective with other kinds of high frequency noise, including "scratch" noise from old 78 rpm recordings, and even certain kinds of electrical buzz. For 78 rpm noises, the noise reduction may work better if the Gate/Exp Attack is increased to 5 msec, with the Gate/Exp Release left at 1 msec.

Setup name: C1 Noise Reducer 2

This is a very powerful noise reducer also designed to reduce hiss and other high frequency noises. The Gate/Exp module is used as a noise reducer identical to that in the above noise reducer. The Comp/Exp section, however, is configured as a second bandsplit low-level expander also capable of achieving noise reduction. The two sections combined give twice the amount of expansion action at low levels. This setup is useful where the maximum degree of noise reduction is required without too much tonal effect on the wanted signal.

The adjustments to C1 Noise Reducer 2 are essentially just three controls:

(i) Click-and-drag on the highlighted button between threshold and GateOpen.

Moving this upward will increase the amount of filtering both of the signal and noise, moving it down will decrease the amount of filtering. Normally, one would use the highest setting that does not audibly degrade the wanted signal or cause objectionable modulation noise. The precise setting depends on the skill and aims of the user.

(ii) Filter frequency.

This adjusts the center frequency of the frequency band subjected to filtering. Normally, this would lie somewhere between 5 and 10 kHz depending on the tonal balance of the hiss, with 7.5 kHz being a typical frequency for "white" hiss noise.

(iii) Filter Q.

This adjusts the width of the filtered frequency band. The lower the Q, the wider the band and greater the amount of perceived noise reduction, but lower Q's also increase the risk of audible modulation noise effects. In practice a Q setting of between 0.5 and 0.6 is usually best for this reason.

Hint: When adjusting the filtering, you may like to switch the Monitor button to Passive mode. This monitors the "passive" band (shown on the frequency response graph in green), which corresponds to the maximally filtered signal. This allows you to hear the effect of the filtering on its own without the dynamic processing, and is useful for tuning for minimum hiss audibility.

Switching between Passive and Audio monitor mode, by clicking on the Monitor button also provides a direct comparison between conventional filtering and the noise reduced audio.

Longer attack times of 5 ms may be used in the two modules - this will have the effect of making the noise reducer 2 work better with impulsive noises such as those on 78 rpm records. Leave the release time at 1 msec.

Either module may be switched out by clicking the Comp/Exp or Gate/Exp button to bypass the relevant module.

Setup name: C1 Gate Noise Reducer

This noise reducer uses a bandsplit gate rather than an expander to cut out noise. This is extremely effective in quiet passages, but results in audible sudden appearance and disappearance of hiss. This can be acceptable when one sound source (e.g. a noisy keyboard or electric guitar) is part of a mix. This setup, also includes, in the Comp/Exp module, a second gentle low-level bandsplit downward expander to provide some noise reduction above the gate threshold, so as to make gate switching effects on noise less audible.

The adjustments for the Gate Noise reducer are the same as for Noise Reducer 2 above.

In addition, The GateClose (normally set between 0 and 9 dB lower than GateOpen) and Hold controls of the gate can be used in the usual way to keep the gate open longer (during which it will let through more noise) if desired to preserve more treble of the wanted sound..

Rumble Reduction

Setup name: C1 Rumble Reducer

This setup reduces bass rumble noises in much the same way that the previous setup reduced hiss noises. Many recordings have excessive low frequency or "rumble" noise, either due to traffic, air conditioning or similar sound sources, or due to air currents or other spurious noises transmitted to the microphone body. Rather than filter out all low frequency sounds to remove these, a rumble reducer that retains bass response for higher level sounds is often preferable, and can have virtually inaudible action.

Its mode of operation is identical to the previous noise reducer with just two exceptions: the filters are now tuned for removing bass frequencies rather than treble, and (ii) the Gate/Exp Attack and Release time constants are set to be longer to minimize bass distortion.

The filter frequency control may be adjusted by ear with the Monitor button in Passive mode to minimize the audible rumble noise, Clicking on the Monitor button back to Audio mode, the GateOpen threshold may then be adjusted to reduce rumble level while having the minimum effect on the tonal quality of wanted sounds.

As before, the Comp/Exp section is set up as a gentle low-level expander to reduce any residual noise in quiet passages, and it may be adjusted by moving the threshold up or down. It may be removed from the algorithm by clicking the Comp/Exp button into bypass mode.

Hint: If the rumble level is very high, it may not be possible to remove rumble without some tonal effect on the wanted signal. In this case, one should first process the sound file to remove the very lowest rumble frequencies by using the Waves Q10 parametric equalizer with a steep cut high-pass filter, for example one taken from the Q10 setups Library bandlimit filters such as the high-pass filter from the Broadcast44.1.setup, adjusted to cut off

at the highest acceptable frequency. In the case of deep rumble, this will remove the largely inaudible but very high energy very low frequencies. The C1 Rumble reducer should then be used on the processed file to reduce the more audible remaining components of rumble.

If a Q10 parametric equalizer is not available, the C1 itself may be used as a high pass-filter for the first stage of processing by loading the setup C1 high-pass filter described below. This may be adjusted in frequency and used for processing to remove the lowest rumble frequencies before the C1 Rumble reducer is applied. However, the sharpness of filtering is less good than can be achieved using the Q10.

Setup name: C1 High-pass filter.

This uses the "Active" monitor mode position (normally used for monitoring the sidechain signal) as a 36 dB per octave high pass filter for applications such as cutting out infrasonic and deep bass rumble sounds. The cut-off frequency may be adjusted by adjusting the frequency Value Window, which actually shows the frequency at which the filter is 9 dB down. An octave higher the filter is only about 1.5 dB down, and an octave lower about 30 dB down. This filter may be used as the first stage in processing where the C1 Rumble Reducer described above is used as the second stage of processing.

Removal of miscellaneous noise

The reduction of noise or background sounds with arbitrary frequency spectrum, rather than just hiss or rumble, involves a greater degree of user patience, experimentation and skill in setting up the C1. The method is as follows.

Load either the C1 Noise Reducer or the C1 Rumble Reducer setup. Initially, click the Monitor button to Passive mode, and tune the filter, using the type, frequency and Q controls for a useful degree of reduction of the unwanted background noises, while making sure that the removal of wanted sounds is not excessively severe. (It is in making this judgement that personal skill and experience are important). Then click the Monitor button back to Audio mode.

One should set up the Attack and Release times of the Gate/Exp section according to the active band response shown in red on the frequency response display. The rule here is to make Gate/Exp Attack and Release times roughly equal, or Release perhaps only twice as large as Attack, and to set Attack at around 1 ms if the active band response shown in red on the frequency response display is shown as below say -15 dB at all frequencies below 1 kHz. If the active band response is more extended at low frequencies, then longer Gate/Exp attack and release times should be used to prevent nonlinear distortion, with Attack and Release times being increased inversely as the bass response frequency of the active band is lowered, up to around 7 ms for an extended bass component in the active band.

One then sets the GateOpen threshold to achieve the desired trade-off between noise reduction and unwanted effects on the wanted signal. In the event that a satisfactory degree of noise reduction is obtained only at the expense of excessive low-level "pumping" of the signal, there are two things you can do to improve the trade-off. First, you can reset the filter so that the active band affects a smaller frequency range, affecting less of the wanted signal. If this does not produce the desired results, you can use the floor control set to somewhere in the range -5 dB to -15 dB to reduce the amount of wanted-signal pumping. This will also reduce the degree of noise reduction, but may give a more satisfactory trade-off, achieving moderate noise reduction but with fewer side effects.

In some situations, you may prefer to set the Gate/Exp module mode to "Gate" rather than "Expander", especially if the noise is only obtrusive in "quiet" gaps in the program material and not during loud passages, and generally in that application you may use somewhat longer Gate/Exp release times of up to say 100 msec. Such gate settings may be used also if the sound is just one component in a mix.

Chapter 7 - De-Essers

On De-Essers

De-essers are essentially dynamic processors that reduce the signal gain during peak-level sibilants ("esses") by operating a gain reduction during high levels in the "Ess" frequency band typically centered at about 7 kHz. They are used to reduce excessive sibilance in speech and other vocal sounds.

Two different kinds of de-essers can be implemented on the C1:

Sidechain de-essers, which reduce the gain of all frequencies in the signal during a high level in the "Ess" frequency band. This is the most familiar type of de-esser, and

Bandsplit de-essers, which reduce the gain only in the "Ess" frequency band but leave other frequencies unaltered.

De-essers also differ in the degree of gain reduction as the "Ess" level increases, and different implementations produce different tonal qualities to de-essing. Which is preferred is partly a question of subjective choice, and partly a question of what other functionality (e.g. compression, gating, noise reduction, etc.) is also available with a particular de-esser in the C1.

Which kind of de-esser?

There are 5 different de-esser modes possible on the C1. These are:

(1) using the Comp/Exp module in sidechain mode

(2) using the Comp/Exp module in split mode

(3) Using the Gate/Exp module in Expander sidechain mode + gain makeup

(4) Using the Gate/Exp module in Expander split mode + Comp/Exp makeup in split mode .

(5) Using the Gate/Exp module in Gate split mode + Comp/Exp makeup in split mode.

Each of these de-essers have a different sound and also allow use with different combinations of processing. So the choice will be made partly on sound and partly on other processing one wishes to do at the same time.

Sidechain v. bandsplit de-essers

Sidechain de-essers have the least tonal effect on the sound, their downside being that their gain variations can sometimes be heard as gain pumping or instability.

Bandsplit de-essers work by varying the gain in the Ess-band. Generally, these have far fewer gain instability effects, since they only affect the Ess band, but their downside is that they can also cause a tonal dulling. This effect tends to be least if the Gate-Split de-esser is used and most if the Comp/Exp module split de-esser is used.

The setups below all use one dynamics module for the main de-essing and the other for other tasks. Where stated, the other module may be switched out if not required.

Compressor + De-Esser

Setup name: C1 Compressor + DeEsser 1

This setup combines a compressor (here set to mid-level compression using the Comp/Exp module in PeakRef & wideband mode) and a de-esser implemented using the Gate/Exp module in sidechain expander mode. The -12 dB output gain is intended to compensate for the +12 dB floor gain used in the de-essing in the Gate/Exp module. Any additional desired change of output gain may be used.

The compressor may be switched in or out by clicking the bypass button above the Comp/Exp module meter bars. The de-esser may be switched in or out by clicking the Gate/Exp bypass button above the Gate/Exp meter bars.

De-Essing adjustments

The GateOpen level is used to adjust the degree of de-essing. You can see by the grey area at the top of the gain reduction meter of the Gate/Exp module, and by the level metering on the light blue line on the input/output graph how the high-level sibilant sounds are being reduced in level. The frequency band of de-essing, shown in red on the frequency response graph, may be adjusted by altering the filter frequency (typically in the range 5000 to 11000 Hz) and the Q (typically 0.4 to 0.6).

Compressor adjustments

The ratio (typically in the range 1.0 to 3.0) of the compressor, and its threshold (typically in the range -35 to -60 dB) may be used to adjust the compression effect to taste, and the attack, release and PDR may also be altered as usual. The degree of low-level gain increase given by the compressor is indicated by the yellow bar of the gain reduction meter of the Comp/Exp module.

De-Ess & De-Hiss

Setup name: C1 De-Ess + De-Hiss 1

This setup uses the Comp/Exp module in sidechain mode as a de-esser, and the Gate/Exp module in split expander mode as a de-Hisser, They share the filter setting, but are otherwise independent. Either module may be bypassed or substituted by another setup for that module without affecting the functioning of the other.

The de-esser may be switched in or out by clicking the bypass button above the Comp/Exp module meter bars. The de-hisser may be switched in or out by clicking the Gate/Exp bypass button above the Gate/Exp meter bars.

De-essing adjustments

The Comp/Exp module is set up to "duck" in level whenever high levels arrive in the active "Ess" frequency band. Altering the threshold control alters the degree of de-essing. You can see by the level metering on the yellow line on the input/output graph and the gain reduction meter (shown in color as a red bar) in the Comp/Exp module how the high-level sibilant sounds are being reduced in level. The ratio control in the range say 2 through 50 to -5 can also be used to alter the quality of the de-essed sound. Generally, ratio settings near 2 or 5 give a moderate degree of de-essing but few unnatural side effects. Ratios near 20 to 50 or -50 to -25 give more effective de-essing, but more side effects, and negative ratios if overdone with too low a threshold can result in an over-processed effect that sounds like the speaker has lost his/her teeth!

De-hissing adjustments

The Gate/Exp module is set up as a de-hisser, and the degree of de-hissing can be adjusted by moving the GateOpen threshold up or down until the effect on tonal quality is minimal. This allows a trade-off between degree of hiss reduction and alteration of tone quality. You can see from the gain reduction meter (shown in color as a red bar) in the Gate/Exp module how much hiss reduction is taking place.

Small adjustments in frequency and Q can be used to optimise either the de-hissing or de-essing, or some compromise between the two.

Hint. When using both effects together, in setting up the de-hissing or de-essing, it is helpful to switch the other processing module to bypass by means of the bypass button above the bar-meter section of the other module. That way one can avoid confusion caused by the interaction of the two processes. Once each is set up individually, switch in both and make final adjustments to each to minimise any tonal dulling. Note that to avoid dulling, the threshold should always be set at least 15 dB, preferably more, below GateOpen.

De-Ess + LowLevel compressor

Setup name: C1 De-Ess + LLcompress

This setup uses the Comp/Exp module in sidechain mode as a de-esser and the Gate/Exp module in wideband expander mode with positive floor gain as a low-level compressor to make sounds louder.

The de-esser may be switched in or out by clicking the bypass button above the Comp/Exp module meter bars. The low-level compressor may be switched in or out by clicking the Gate/Exp bypass button above the Gate/Exp meter bars.

The Comp/Exp module is set up to "duck" in level whenever high levels arrive in the active band. Altering the threshold control alters the degree of de-essing. The ratio control in the range say 2 through 50 to -5 can also be used to alter the quality of the de-essed sound.

The filter may be tuned in frequency and Q slightly to alter the predominant de-essing frequency band. One can monitor the selection of Ess frequencies by switching the monitor button to "sidechain".

Control of low-level compression

The floor control, between 0 and 12 dB and the GateOpen threshold of the expander module can be used to adjust the degree of low-level compression. Higher settings of floor give a higher degree of compression. The threshold adjusts the sound level below which compression occurs. The attack and release times of the low-level compressor can be adjusted to taste in the usual way for a compressor. Use of lookahead in the "yes" position is usually recommended.

Bandsplit De-Ess, De-Hiss and EQ

Setup name: C1 De-Ess + EQ + De-Hiss 2

This is very similar to C1 De-Ess + De-Hiss 1 except that now the de-esser is working in bandsplit rather than sidechain mode. Adjustments are made in much the same way as for the C1 De-Ess + De-Hiss 1 setup, although the sound quality of this de-esser is rather different, often requiring significantly lower threshold settings.

A second difference of this setup is that it can additionally be used to alter the EQ by altering the level of the active band, using the makeup gain to boost or cut the active band. This helps to compensate for any slight dulling of mid-level sounds given by the combination of de-essing and de-hissing, or may be used to provide additional tonal alteration.

Bandsplit De-Hiss, De-Ess and EQ

Setup name: C1 De-Hiss + EQ + De-Ess 3

This setup combines a de-hisser using the Comp/Exp module in PeakRef & bandsplit mode) and a de-esser implemented using the Gate/Exp module in bandsplit expander mode. The -12 dB makeup gain in the Comp/Exp module is intended to compensate for the +12 dB floor gain used in the de-essing in the Gate/Exp module. Any additional desired change of output gain may be used.

The de-hisser may be switched in or out only by altering the value of ratio to 1, since its make-up gain is part of the functioning of the de-esser. The de-esser may be switched out by clicking the Gate/Exp bypass button above the Gate/Exp meter bars and also altering the makeup gain to 0 dB.

De-Essing adjustments

The GateOpen level is used to adjust the degree of de-essing. You can see by the grey area at the top of the gain reduction meter of the Gate/Exp module, and by the level metering on the light blue line on the input/output graph how the high-level sibilant sounds are being reduced in level. The frequency band of de-essing, shown in red on the frequency response graph, may be adjusted by altering the filter frequency (typically in the range 5000 to 11000 Hz) and the Q (typically 0.4 to 0.6).

De-Hisser adjustments

The threshold (typically in the range -65 to -100 dB) may be used to adjust the de-Hissing effect to obtain the best compromise between hiss level and alteration of tonal quality of the wanted sound. The degree of hiss reduction at each moment is indicated by the red bar of the gain reduction meter of the Comp/Exp module.

You should try and ensure that the de-hissing threshold is not too close to the de-essing threshold, since otherwise there will be a significant loss of the "ess" frequencies at all signal levels.

EQ adjustments.

You can alter the overall tonal quality further by adding additional EQ in the active band. This may be done by varying the Makeup gain up or down around its normal level. This can also help to compensate for tonal losses due to the combined effect of the de-hissing and de-essing.

Setup name: C1 De-Hiss + EQ + De-Ess 4

This is used identically to the previous set-up, differing from it only in that the Gate/Exp module is now in gate rather than expander mode.

However, the more extreme sudden change of gain above the "Ess" threshold given by the gate means that the subjective effect of this setup is very different. In some cases, it will be found to be very effective at de-essing, but in other cases, the effect may be found to be overdone.

Chapter 8 - Enhancers

About enhancement

Quite often one wishes to add clarity and/or depth to a soundfile to give it more "pizazz", intelligibility or impact. This can be done using conventional equalization, but all too often this has the side effects of making bass boomy or treble harsh and uncomfortable. The C1 enhancement modes allow clarity and depth to be added either to a single sound or to a complete mix with far fewer unpleasant side effects.

The general idea of these enhancers is to boost bass or treble frequencies only at lower levels below a threshold, but not at high levels. This ensures that the loudest, and potentially most oppressive, parts of the sound do not have an exaggerated bass or treble, while the quieter parts have the extra warmth or sparkle obtained by being boosted in level. Generally, the choice of boosted frequencies is fairly critical for optimum effect, although the choice may depend to some extent on the sound being processed.

Speech and vocal enhancer (with optional mid-level compressor).

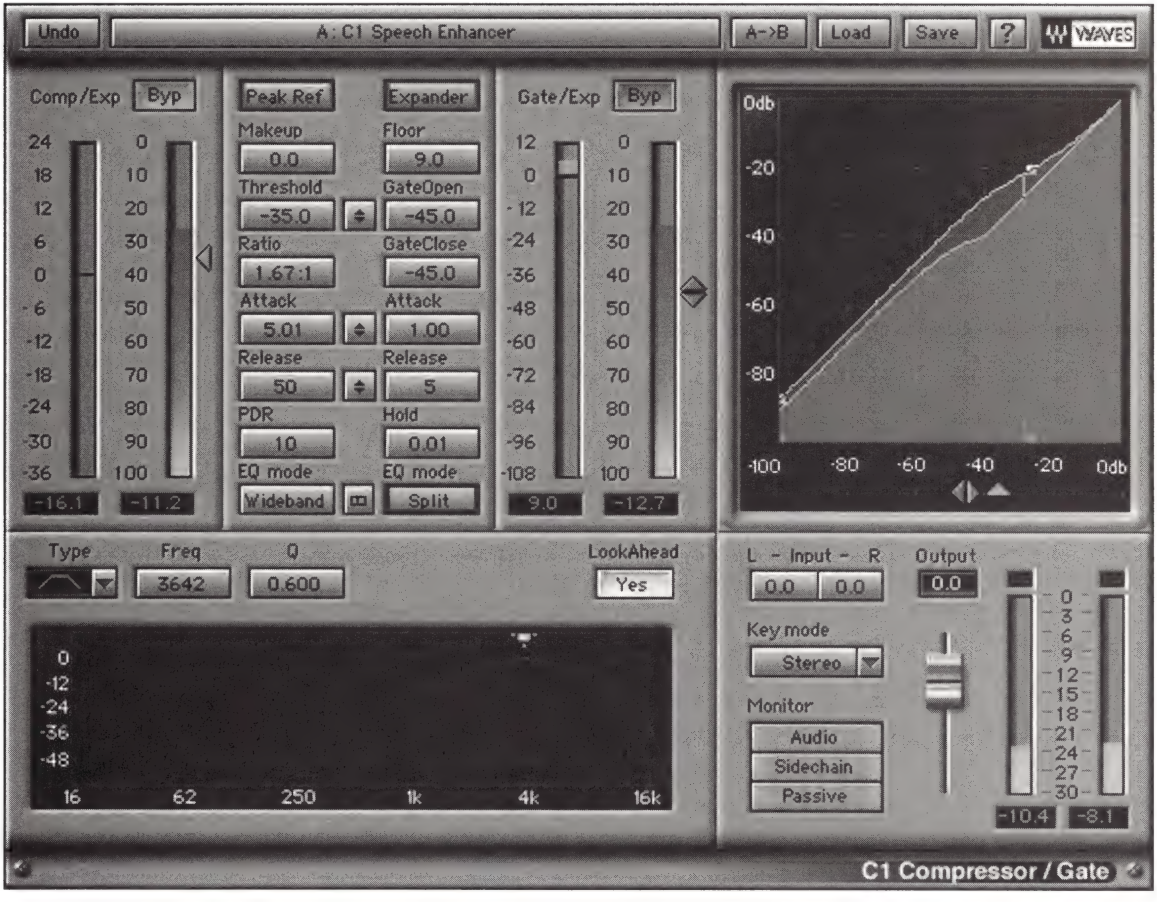

Setup name: C1 Speech Enhancer

The Gate/Exp module is here used in bandsplit mode to boost speech presence frequencies for enhanced intelligibility without the degree of unpleasant harshness or sibilance associated with conventional presence EQ. These frequencies are boosted by the setting of the floor control at low levels only, the threshold for those levels being set by the GateOpen threshold control.

The selective boosting only of lower levels in the presence band achieves a clarity and an ability for speech or vocals to cut through noisy listening environments (or dense mixes) without the unpleasantness given by simple EQ used to the same end, and the results still sound reasonable on high-quality reproduction systems.

Adjustment

The degree and subtlety of the effect may be varied by adjustment of the floor level boost between 0 and +12 dB (which acts as a presence gain control), with higher settings giving a greater effect, and the GateOpen threshold - lower GateOpen settings of threshold give a more subtle effect and higher ones a more blatant effect.

This is an effect that can be used subtly, almost subliminally, to just make speech (or musical vocals) seem more "present" without any very obvious change of tonal quality in a mix, or it can be used to provide blatant tonal changes with fewer than usual side effects.

This is usually best with frequency and Q settings roughly as shown - lower frequencies tends to start emphasising harshness, and higher frequencies emphasise brightness rather than vocal intelligibility - however, this may be an effect that may be wanted to brighten mixes without serious side effects.

The Comp/Exp module , shown bypassed above, is set up for a moderate degree of wideband mid-level compression, suitable for making speech louder without making peak levels drive the reproduction system too hard too much of the time. The compression effect may be adjusted by moving the threshold up or down and the compression ratio between say 1 and 3.

Brightness enhancer + compensating noise control.

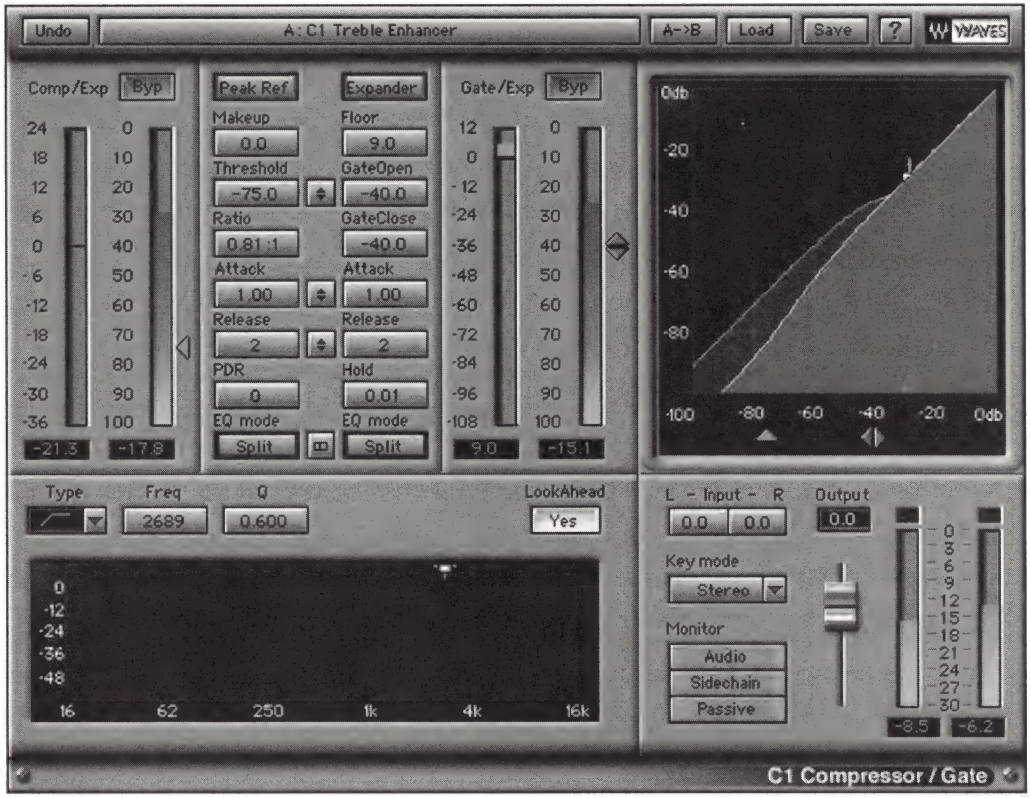

Setup name: C1 Treble Enhancer

The primary function of this setup is to act as an effective "brightness enhancer". This is provided by the Gate/Exp module in split mode. It will be noted that many consumers prefer the sound of Dolby™ B cassettes without Dolby™ decoding because of the low-level brightness enhancement done by that process, and the basic enhancement idea here is similar, but more controllable. High frequencies are boosted by up to 9 dB, the degree of boost being controlled by the floor control, and this boost takes place below a threshold set by the GateOpen threshold. These two controls may be tuned - floor between 0 and + 12 dB and GateOpen typically between -20 and -60 dB, for optimum enhancement of a particular mix or sound. The effect gets more subtle as GateOpen and Floor are lowered, and more blatant as they are raised.

The yellow bar on the Gate/Exp gain reduction meter can be used to monitor the degree of gain increase in the treble moment by moment.

The Comp/Exp module is here used in an unusual low-level expansion bandsplit mode to solve the problem that the enhancer also raises the level of any low-level hiss noise present, by acting as a gentle noise reducer at very low levels to pull the noise back down again. Generally speaking, the threshold should be set very low so as to ensure

that the noise reduction does not affect the mid-levels that are being boosted by the enhancer. The threshold should be adjusted by ear to ensure that the noise reduction is effective without affecting enhancement effects too much.

The red bar on the Comp/Exp gain reduction meter can be used to monitor the amount of low-level treble noise reduction moment by moment.

There is nothing sacred about the filter frequency and Q settings used in this setup - but it is often found that the high-pass setting works best at maximum Q of 0.6. The setting shown is one that generally increases brightness without too much mid range harshness creeping in. You may also care to experiment using settings with bandpass filter type tuned with frequency in the region 3 to 9 kHz and having Q's in the region 0.3 to 0.6.

Bass enhancer + compensating rumble control.

Setup name: C1 Bass Enhancer

This is the bass equivalent of the brightness enhancer above.

The primary function of this setup is to act as an effective "bass enhancer". This is provided by the Gate/Exp module in split mode. It provides an added bass richness or depth, even played back via modest equipment, without making the bass overwhelming and unpleasant. Selected low frequencies are boosted by up to 9 dB, controlled by the floor control, and this boost takes place below a threshold set by the GateOpen threshold. These two controls may be tuned - floor between 0 and + 12 dB and GateOpen typically between -15 and -50 dB, for optimum enhancement of a particular mix or sound. The effect gets more subtle as GateOpen and Floor are lowered, and more blatant as they are raised.

The yellow bar on the Gate/Exp gain reduction meter can be used to monitor the degree of gain increase in the bass moment by moment.

The Comp/Exp module is here used in an unusual low-level expansion bandsplit mode to solve the problem that the enhancer also raises the level of any low-level rumble noise present, by acting as a gentle noise reducer at very low levels to pull the noise back down again. Generally speaking, the threshold should be set very low so as to ensure that the noise reduction does not affect the mid-levels that are being boosted by the enhancer. The threshold should be adjusted by ear to ensure that the noise reduction is effective without affecting enhancement effects too much.

There is nothing sacred about the filter settings used in this setup - but it is often found that the bandpass filter type setting with frequency tuned to predominant frequencies in the range 100 to 150 Hz works best at maximum Q of 0.6. The setting shown is one that generally increases the bass in contemporary popular music without too much lower mid range muddiness creeping in. You may also care to experiment using lowpass filter type setting with frequency tuned in the frequency region up to 350 Hz and having Q's of 0.6.

Generally, the lower the bass frequencies selected, the longer should be the Gate/Exp release times to minimise possible bass distortion effects.

Bass & Treble enhancer + compensating noise control.

Setup name: C1 Bass/Treble Enhancer

This is the equivalent of the two enhancers above when enhancement of both frequency extremes at the same time is required.

The primary function of this setup is to act as an effective "bass and treble enhancer". This is provided by the Gate/Exp module in split mode. It provides an added bass richness or depth and an added brightness, even played back via modest equipment, without making the bass or treble overwhelming and unpleasant. Selected low and high frequencies are boosted by up to 9 dB, with boost controlled by the floor control, and this boost takes place below a threshold set by the GateOpen threshold. These two controls may be tuned - floor between 0 and + 12 dB and GateOpen between -15 and -50 dB, for optimum enhancement of a particular mix or sound. The effect gets more subtle as GateOpen and Floor are lowered, and more blatant as they are raised.

The yellow bar on the Gate/Exp gain reduction meter can be used to monitor the degree of gain increase in the bass and treble moment by moment.

The Comp/Exp module is here used in an unusual low-level expansion bandsplit mode to solve the problem that the enhancer also raises the level of any low-level rumble and hiss noise present, by acting as a gentle noise reducer at very low levels to pull the noise back down again. Generally speaking, the threshold should be set very low so as to ensure that the noise reduction does not affect the mid-levels that are being boosted by the enhancer. The threshold should be adjusted by ear to ensure that the noise reduction is effective without affecting enhancement effects too much.

There is nothing sacred about the filter settings used in this setup - but it is often found that the band-reject filter type setting with frequency tuned to predominant frequencies in the range 500 to 1000 Hz works best at Q between 0.15 and 0.3 so as to give the desired bass and treble crossover frequencies on the graph. The setting shown with crossover frequencies around 125 Hz and 4 kHz is one that generally increases the bass and brightness in contemporary popular music without too much lower mid range muddiness or upper mid range harshness creeping in.

Cheap System enhancer + compensating noise control.

Setup name: C1 CheapSystem Enhancer

This is a variation of the Bass & Treble enhancer above, designed to solve the problem of an enhanced sound designed to be played on cheap crummy systems, such as many multimedia systems, AM radios, cheap portable headphone systems or even utilitarian PA and muzak-type systems, preferably while still sounding good on high-end systems. It actually manages to incorporate three functions: a high level bandwidth reducer, a low level bass-and-treble enhancer, and a single ended noise reducer.

Cheap systems often have limited frequency range at the two frequency extremes, and they often distort if one attempts to push much level into them at the frequency extremes, so that limitation of bandwidth at high input level is necessary to get the best out of such systems - but this often makes the sound terrible over better systems.

This setup combines bandwidth limitation between 250 and 4000 Hz at high levels - reducing levels at the frequency extremes by 6 dB, with 12 dB bass and treble enhancement at low levels. This subjectively causes an overall increase in bass and treble - even on cheap systems, with a nice-sounding enhancement, while actually limiting the degree to which cheap systems are exposed to signals they cannot handle. The result is that cheap systems can be played louder without distortion and yet give enhanced bass and treble at the same time. The effect is sufficiently subtle that the enhancement also sounds good on high-quality systems.

The main control here is the GateOpen threshold control, which should be set somewhere between -15 and -35 dB on typical programme material. The gain Makeup control sets the degree of high-level attenuation of bass and treble, which should always be less than the degree of boost of floor, which sets the amount of lower-level enhancement. The Comp/Exp threshold control, which should be set low by ear to have minimum side effects on the wanted signal, acts as a low-level noise reducer for any noise present in the source material.

The yellow bar on the Gate/Exp gain reduction meter can be used to monitor the degree of gain increase in the bass and treble moment by moment, but the degree of gain increase in this setup is reduced by the makeup value, in the case shown above, by 6 dB.

Compressors + Enhancers

These 4 setups use the Comp/Exp module in wideband mode for compression and the Gate/Exp module in bandsplit mode for enhancement. The two modules operate effectively independently on one another. These setups are particularly useful for maximising impact of sound files that have to be squeezef into a narrow dynamic range, such as for multimedia or broadcast material.

Setup names:
- C1 HLcompress+bass/treb enhance
- C1 HLcompress+treble enhance
- C1 MLcompress+bass/treb enhance
- C1 MLcompress+treble enhance

The Gate/Exp and Filter modules of these setups are the same, and used in the same way, as those for the C1 bass/treble enhancer or the C1 treble enhancer above.

The Comp/Exp module of these setups is used as a compressor. The HLcompress setups are high-level compressors that may be used conventionally, similarly to as describedefor the C1 HLcompress + De-Hiss setup. The MLcompress setups are mid-level compressors similar to, and used in the same way as, the Comp/Exp sections of Chapter 4: C1 Multimedia Speech 1, C1 Multimedia Speech 2 and Chapter 5: C1 Compressor + De-Reverb.

Enhancers + Gate

These two setups provide enhancement plus gate, using the Comp/Exp with Filter modules as an enhancer and the Gate/Exp module in Wideband Gate mode as a gate.

Setup name: C1 Treble enhance + Gate

The Comp/Exp module in split mode of this setup acts as an effective "brightness enhancer". Low-level high frequencies are boosted by up to 10 dB, the degree of boost being controlled by the ratio control, and this boost takes place below a threshold set by the threshold. These two controls may be tuned - ratio between 1 and 1.60 and threshold typically between -40 and -80 dB, for optimum enhancement of a particular mix or sound. The effect gets more subtle as threshold and ratio are lowered, and more blatant as they are raised.

The yellow bar on the Comp/Exp gain reduction meter can be used to monitor the degree of gain increase in the treble moment by moment.

There is nothing sacred about the filter frequency and Q settings used in this setup - but it is often found that the high-pass setting works best at maximum Q of 0.6. The setting shown is one that generally increases brightness without too much mid range harshness creeping in. You may also care to experiment using settings with bandpass filter type tuned with frequency in the region 3 to 9 kHz and having Q's in the region 0.3 to 0.6.

The Gate is a conventional fully-functioned wideband gate. The GateClose threshold is preset 3 dB lower than the GateOpen threshold so as to reduce gate "chatter", and the hold control ensures a minimum gate on time which again helps prevent chatter. Attack and Release may be adjusted in the usual way. If it is desired to retain some of the background atmosphere during quiet passages, the floor control may be used to adjust the remaining signal level when the gate is "off".

Setup name: C1 Bass/treb enhance+Gate

The Comp/Exp module in split mode of this setup acts as an effective "enhancer" of ther bas and treble frequency extremes. Low-level low and high frequencies are boosted by up to 10 dB, the degree of boost being controlled by the ratio control, and this boost takes place below a threshold set by the threshold. These two controls may be tuned - ratio between 1 and 1.60 and threshold typically between -40 and -80 dB, for optimum enhancement of a particular mix or sound. The effect gets more subtle as threshold and ratio are lowered, and more blatant as they are raised.

The yellow bar on the Comp/Exp gain reduction meter can be used to monitor the degree of gain increase in the bass and treble moment by moment.

There is nothing sacred about the filter frequency and Q settings used in this setup - and the bass and treble ranges that are enhanced may be adjusted by varying the Q and frequency around the preset values.

The gate is adjusted in the same way as in the previous setup.

Chapter 9 - Keying Setups

About keying

These setups make use of the key mode button, and allow a signal in one of the two channels to control the dynamics of the other.

The most common uses of this kind of facility is ducking and gating. Ducking reduces the gain of one signal when the other is above a threshold. It is normally used to make a background sound quieter when a foreground sound must be heard over the background. It is commonly used with a speech signal to control the dynamics of loud background noises such as audience applause during a broadcast announcement or commentary, so that the speech is clearly audible over the background.

The other common use is keyed or triggered gating, where the level in one track is used to operate a gate in a second track, so that sounds in the second track are allowed through only rhythmically in time with the occurrence of sounds in the first track. For example, one can use a drumbeat on the left track to control the gating of a steady sound in the right track so as to create a different percussion sound having the tonal quality of the steady sound, in time with the original drumbeat. By using sidechain filtering, one can pick out particular frequency components of the drumbeat such as a snare, bass drum or cymbal sound to key or trigger the other sound.

By using bandsplit processing with ducking or triggered gating, one can process just a selected frequency band rather than the full signal spectrum.

Keying can also be used with expanders instead of gates, which has particular uses in creating rhythmic tracks out of steady sounds such as natural steady sound effects like wind noises, rainfall, babbling streams, waterfalls, crowd noise, audience applause etc., as well as artificial continuous sounds like looped sounds, synthesiser tracks etc. If used with bandsplit, this can give otherwise steady sounds that vary in quality rhythmically in time with the original keying track. A sound effect can be made to take on musical rhythmic qualities by being transformed by a keying effect. This leads to a considerable variety of creative musical effects in dance and ambient musics, as well as ways of integrating sound effects and music in radio, TV, video and film drama, documentary, games and advertising applications. The keyed EQ expansion effect described below is particularly flexible and effective in these applications.

Preparing and processing files for keyed effects

Because Plug-Ins for Sound Designer II™ are stereo processors capable of taking in and processing only two channels of sound at a time, to obtain the keying effect, it is first necessary to prepare a sound file having the keying sound that controls the dynamics on one channel, and the sound that is to be dynamically processed - the keyed or triggered sound on the other. The C1 provides the option for either channel to be allocated to the keyed or keying sound. For ease of description, we assume that the keying or triggering sound is in the left channel and that the sound to be dynamically processed or keyed is in the right channel. For this option, the Key Mode button is in the L -> R mode. All keying setups are loaded in this mode.

When processing sound files in this mode, only the dynamically processed sound on the right channel will be present, with nothing on the left channel. If you wish to preserve the original keying sound as well on the left

channel, in an unprocessed form, you should before processing select the right channel signal only. Processing will then result in a stereo file with the unprocessed keying signal on the left, and the processed keyed signal on the right channel. This is useful in cases such as ducking where it is usually desired to mix these two signals together at a later stage.

Hint: If you wish to process or key a stereo signal by a mono keying signal, this cannot be done directly on the C1 because of its limitation to two channels of input. However, a stereo workaround exists as follows. Suppose that you start off with a file K of the keying signal and a stereo file B of the stereo signal to be processed. Use the mix option of Sound Designer II to mix two new stereo files BL and BR, where BL has a mono mix of K on the left channel and the left channel of B on the right channel of BL, and where BR has a mono mix of K on the left channel and the right channel of B on the right channel of BR.

In order to preserve stereo balance later, use 100% gains for transferring channels of B to BL and BR, DO NOT use scaling, and ensure that all time delays in the mixing are zero. Then process both BL and BR with the setup you have chosen, taking care not to change any control values when processing the two files (again to ensure stereo balance is preserved). The resulting processed files PL and PR may now again be mixed into a new stereo file P by putting the right channel of PL onto the left channel of P and the right channel of PR onto the right channel of P again using the Sound Designer II mix option. P will be the desired processed stereo file.

If one needs to mix the processed file with the unprocessed keying signal K, the last stage of preparing P may be bypassed, and the final mixed file M be obtained by mixing from PL, PR and K, putting the right channel of PL onto the left channel of M, the right channel of PR onto the right channel of M, and the original sound file K mixed and panned into M as desired. The C1 introduces no relative time delays between files even when lookahead is used, so that the files will stay in synchronism.

If the keying sound is on the right channel and the sound to be dynamically processed is on the left channel, click the Key Mode button to the R -> L position. Where relevant, select the left channel signal before processing in this case.

Ducking

Ducking is an effect in which the gain of the keyed sound is reduced when the keying sound is louder than a user-determined threshold level. It is typically used to mix speech and background sounds where the background sound would normally drown out the speech, e.g. applause. If the speech or vocal sounds are used as the keying signal, and the applause signal gain is reduced in response to the speech, then the applause will be reduced in level enough that the speech will remain audible when the unprocessed speech is mixed with the ducked applause, because the applause level will be reduced during speech. Many similar applications exist where an important sound can be used to reduce the level of less important sounds to make it more audible.

For example, in recordings of amplified music, a mix from direct feeds from microphones and instruments can be used to duck ambience microphones, so that audience response is audible when the music is not playing, but the sound of the ambience microphone is reduced so as not to muddy the sound of the mix when the music is playing.

When using ducking processing, it is a good idea only to process the right channel (when using L -> R key mode), so that the processed file has the unprocessed keying signal (e.g. speech or music) on its left channel and

the ducked signal (e.g. applause) on its right channel. The mix can then be done direct from the processed file.

The setups library provided two kinds of ducking. The first continues to duck levels further as the level of the keying sound (e.g. speech) increases further, whereas the second gives a fixed user-determined gain reduction when the keying sound is at any level above threshold. The second type of ducking may be preferred as giving a more natural result, but it is very much up to the user which is considered most appropriate.

Setup name: C1 Ducking 1

Here, the Comp/Exp module is used as a ducking device. The peaks of the sound on the left channel cause the right channel to reduce during those peaks.

The threshold level of the keying signal at which gain reduction starts is controlled by the Threshold control. The ratio control, which typically may be set between 2 and 50 or -50 and -15, controls the degree of ducking above threshold. The time duration for which the gain reduction lasts is controlled by the Comp/Exp Release control. The longer this is set, the fewer disturbing rapid level fluctuations will be heard in the ducked sound.

An added refinement is the use of the sidechain mode, so that it is the peaks in the equalized left channel sound that causes the right channel to duck. The above is actually tuned so that the right channel ducks during speech presence peaks on the left channel.

The active frequencies of the keying signal that operate the gain ducking may be selected by operating the filter type, frequency and Q controls. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position. Keying by the wideband signal may also be selected by clicking the Comp/Exp EQ mode button to Wideband.

Setup name: C1 Ducking 2

Here, the Gate/Exp module is used as a ducking device. The peaks of the sound on the left channel cause the right channel to reduce during those peaks.

The threshold level of the keying signal at which gain reduction starts is controlled by the GateOpen control. The floor control, which may be set between 0 and +12 dB, (set to 12 dB in the above example) controls the ducking gain reduction above threshold. The output gain must be set at an equal and opposite compensating gain reduction (-12 dB in the above example) to ensure no gain when the keying signal is low in level. The output gain may, of course, be additionally varied about this value according to the user's requirements. The time duration for which the gain reduction lasts is controlled by the Gate/Exp Release control. The longer this is set, the fewer disturbing rapid level fluctuations will be heard in the ducked sound.

An added refinement is the use of the sidechain mode, so that it is the peaks in the equalized left channel sound that causes the right channel to duck. The above is actually tuned so that the right channel ducks during speech presence peaks on the left channel.

The active frequencies of the keying signal that operate the gain ducking may be selected by operating the filter type, frequency and Q controls. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position. Keying by the wideband signal may also be selected by clicking the Gate/Exp EQ mode button to Wideband.

Ducked EQ

Ducked EQ is a variation of standard ducking, using the C1's bandsplit mode, where the gain reduction is confined to a limited frequency band. This is an extremely useful effect or tool whenever a mix suffers from the common problem of several sounds in a mix all occupying the same frequency range, so making sounds unintelligible. This setup creates "space" in the mix allowing a chosen predominant sound to become audible without losing the other sounds.

It may also be used in situations where conventional ducking is used (e.g. of applause when mixed with speech), where the ducking processing can be confined to a limited frequency band for reduced audible side effects.

Put the sound you wish to predominate in the mix on the left channel of a sound file and one or more of the clashing sounds on the right. Then using L -> R key mode and bandsplit mode for the ducking, ducking is confined to the active band of a filter tuned to the frequency band in which the aural clash occurs. Peaks in the intended predominant sound causes the level of the other sound to "duck" during peaks - but only in the problem frequency band. This minimizes the audible side effects of ducking, since the complementary passive frequency band remains unprocessed.

The interfering sounds duck in level only at the moments they need to. A typical example is a mix with saxes, electric guitar and vocals, all of which clash around 2 kHz. Putting the vocals on the left of the sound file and the sax or guitar on the right ducks the problem frequency range when needed to make the vocals clearly audible, while letting through these instruments the rest of the time and in frequency bands not needed for vocal audibility.

As in the straightforward ducking case, there are two kinds of ducked EQ setup provided in the C1 setups Library. The first continues to duck levels in the active frequency band further as the level of the keying sound (e.g. vocals) increases further, whereas the second gives a fixed user-determined gain reduction in the active frequency band when the keying sound is at any level above threshold. The second type of ducking may be preferred as giving a more natural result, but it is very much up to the user which is considered most appropriate.

Setup name: C1 Ducked EQ 1

Apart from the filter controls, this is used and adjusted in an identical manner to the C1 Ducking 1 setup described above.

The active frequency band in which the ducking operates is shown in red on the frequency response graph, and may be selected by operating the filter type, frequency and Q controls. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position.

The ducked EQ may be bypassed by clicking on the Comp/Exp button immediately above the Comp/Exp bar meters.

Setup name: C1 Ducked EQ 2

Apart from the filter controls, this is used and adjusted in an identical manner to the C1 Ducking 2 setup described above, with one exception: The floor control, which may be set between 0 and +12 dB, (set to 12 dB in

the above example) controls the ducking gain reduction above threshold as before, but the makeup gain must be set at an equal and opposite compensating gain reduction (-12 dB in the above example) to ensure no gain in the active frequency band when the keying signal is low in level. .

The active frequency band in which the ducking operates is shown in red on the frequency response graph, and may be selected by operating the filter type, frequency and Q controls. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position.

The ducked EQ may be bypassed by clicking on both the Comp/Exp and Gate/Exp buttons immediately above the module bar meters.

Keyed Gating

Setup name: C1 Keyed Gate

Here the Gate/Exp module is used as a gate for the signal on the right channel, and the gating is keyed or triggered (i.e. opened) by the signal on the left channel. The gating is tuned by the sidechain filter and the GateOpen and GateClose threshold. The signal on the right channel is switched in time with the beat on the left channel. The longish release time shown here is to get a sense of "decay" after the beat, and the hold time is used to establish a minimum time for which the signal is switched fully on.

The GateOpen threshold may be moved up or down to adjust the triggering of the gate. Gate/Exp Release and Hold times can be varied to alter the "decay" of the percussion envelope effect.

The type, frequency and Q of the sidechain filter can be tuned so that the frequencies of the signal used for keying can be selected. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position. Keying by the wideband signal may also be selected by clicking the Gate/Exp EQ mode button to Wideband.

The keying can be bypassed by clicking on the Gate/Exp button immediately above the Gate/Exp bar meters.

Keyed expander

Setup name: C1 Keyed Expander

This setup is similar to a keyed gate, except that it replaces the gate by a low level (downward) expander. This has the effect of not cutting off the keyed sound during quiet passages, but of modulating its level with the envelope of the keying or trigger sound. This is particularly good for using keying to create a "new" percussion sound from a steady state sound like hiss or a synthesiser drone or even a natural sound effect like a babbling brook (be creative!). Unlike using a gate, a keyed expander has a more subtle and "organically natural" quality, due to its preservation of dynamic variations in the original keying signal.

Both the Gate/Exp and the Comp/Exp modules are set up as almost-identical keyed expanders in this setup, and either may be used for this effect on their own. But if both are switched in, the effect is more dramatic and effective.

The thresholds of either module (i.e. threshold and GateOpen) may be moved up or down together, by click-and dragging on the "button" between the Threshold and GateOpen Value Windows, to vary the threshold level below which expansion occurs. Release times can be jointly varied, by click-and-dragging on the "button" between the two Release Value Windows, to alter the "decay" of the percussion envelope effect.

The type, frequency and Q of the sidechain filter can be tuned so that the frequencies of the signal used for keying can be selected. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position. Keying by the wideband signal may also be selected by clicking the Gate/ Exp EQ mode button to Wideband.

The keying of either or both of the two modules can be bypassed by clicking either or both of the bypass buttons immediately above the module bar meters.

Keyed EQ expansion

This is an especially flexible creative effect, in which a keying sound such as a percussion or drum track alters the equalization of a steady or continuous sound so as to impose the rhythmic percussion effect on the new sound. The flexibility lies is the ability both to tune the frequencies, and hence the tonal quality, of the sound that are effected and the exact degree, from almost subliminally audible to blatant, to which the EQ is varied in time with the original percussion. Additionally, the overall EQ can be varied in this setup.

This effect can be used to add varied rhythmic effects to continuous, drone or natural sounds that may not even be musical in nature, and can be useful in both musical applications and for musical dramatic effects in drama, documentary and other applications. In musical applications, the EQ can be tuned to ensure maximum audibility of the rhythmic effect of the keyed sound within a mix.

Setup name: C1 Keyed EQ Expander.

This setup is based on a keyed expander in bandsplit mode

The degree of boost of the selected active frequency band above the keying threshold is set by the makeup gain and the floor control, which may be set between +12 and -40 dB, must be set to an equal and opposite gain to ensure no gain in the active frequency band when the keying signal is low in level.

Any other degree of EQ in the active band above threshold can be used by setting makeup and output gains to equal and opposite settings. Additional overall adjustment of EQ in the active band at both high and low levels of the keying signal may be provided by varying the makeup gain above or below this setting.

The active frequency band in which the keyed EQ operates is shown in red on the frequency response graph, and may be selected by operating the filter type, frequency and Q controls. The effect of the equalisation on the keying signal may be auditioned by clicking the monitor button to Sidechain position.

The longish release time shown here is to get a sense of "decay" after the beat, and adds a kind of reverberant decay effect to the keyed sound.

The GateOpen threshold may be moved up or down to adjust the keying threshold. Gate/Exp Release time can be varied to alter the "decay" of the percussion envelope effect.

The keyed EQ expansion may be bypassed by clicking on both the Comp/Exp and Gate/Exp buttons immediately above the module bar meters.

Acknowledgements

These setups were compiled using a wide variety of sound materials, including commercial recordings. Of particular use were unprocessed recordings from the Alan Parsons/Stephen Court "Sound Check" test CD, the EBU SQAM (Sound Quality Assessment Material) CD, and a limited edition CD-R "Higher Fidelity from 78s" of careful unprocessed 78 rpm transcriptions by Peter Adamson. Thanks are due to Peter Adamson, Peter Craven, Jerry Bruck, Geoffrey Barton and many others who have contributed by discussions over the years to the know-how incorporated into these setups. Also to Meir Shashoua and Seva for almost superhuman endurance far, far beyond the call of duty in making this all possible.

The text of this setups manual was written by Michael Gerzon, with contributions and Editing by Seva.

S1
Stereo Imager

Table of Contents

Register Now! ..... see WaveSystem Manual

Chapter 1 ..... About the S1 Stereo Imager 3

Before you start 3
Features and applications 3
S1 component plug-ins 4

Chapter 2 ..... The S1 Interface 5

Stereo Vector Display™ 5
Quick guide screen shot 5
The Controls 6
Gain, Width 6
Asymmetry, Rotation, Shuffling, Frequency 6
LR/MS input, Bass Trim, Reset 7
Signal path (Order of action of controls) 7

Chapter 3 ..... Using the S1 8

Graphical control 8
Hints for Use 8
Mastering and general use 9
To mix multiple stereo sources with S1 10
Mono Compatibility 10

Chapter 4 ..... About MS and MS modes .. 11

What is MS? 11
How is MS used? 12
Polarity/Channel swap in MS mode 12
Use of LR/MS meter mode 13

Chapter 5 ..... Other S1 controls .. 14

Peak counters 14
Output clipping (the Gain control) 14

Chapter 6 ..... Special MS uses .. 15

Using the S1 MS options 15
MS Equalization 15
Uses of MS equalization 16
How to do:
Overall EQ 16
Frequency-dependent width 16
One-point microphone compensation 17
FM broadcast noise reduction 17
One-sided equalization 17

Chapter 7 ..... More about shuffling .. 19

Chapter 8 ..... History of stereo controls .. 20

Chapter 9 . What kind of stereo will the S1 work on? .. 21

Chapter 10 ..... Key commands .. 22

Chapter 11 ..... Stereo monitoring .. 23

Chapter 1 - About the S1 Stereo Imager

The S1-Stereo Imager combines a number of classical stereo processing techniques with Waves' intuitive user interface to provide a powerful and unique 'stereo toolkit' for use in mastering and digital editing situations. S1 may be used to correct faults inherent in a stereo recording or mix, or it may be used to creatively enhance the stereoism of existing stereo material.

S1 is based on recognized engineering principles and does not purport to create a three-dimensional soundfield from a two-loudspeaker system or to create a pseudostereo master from a mono source. Because S1 is based on established audio engineering principles, it offers a high level of mono compatibility with minimal side effects.

S1 provides the means to readjust the stereo level-balance of a mix and has the ability to dramatically widen an existing stereo image, again without introducing significant side effects. For example, an assymetrical mix can be adjusted so as to re-balance the left and right components without affecting the position of the centre image or, alternatively, the center image may be panned left or right without affecting the balance of the stereo components of the mix.

Like all Waves plug-ins, S1 is based around a graphical interface and is designed to be both straightforward and intuitive to use. You can adjust all parameters via traditional numerical displays or directly from the graph, giving you the choice of interface and control.

Before you use S1, you should have a good working knowledge of both the host application software and your computer. If you've used a Waves plug-in before, you'll probably be able to find your way around the S1 interface with minimal reference to this manual. However, we recommend you read it through at least once just so you know all the things S1 can do.

Before you start

Please make sure that your monitoring system is in phase and set up symmetrically about your normal listening position! This may seem obvious, but because S1 is designed to operate on the spatial aspects of a mix, it is imperative that your monitor system conveys an accurate sense of stereo imaging. If you have any concerns about your monitoring arrangement, please read the **Stereo Monitoring** chapter.

Features and applications

Particular care has been taken in the design of S1 to avoid unpleasant and fatiguing 'phasiness' effects and to retain a high level of mono compatibility. Mono sounds remain uncolored; S1 does not change the basic tonal character of the original sound, but simply sweetens and broadens the spaciousness of the stereo image.

A classic phase-compensated 'shuffler' is provided to both enhance the stereo spaciousness and improve the image quality at the bass end of the audio spectrum, again without introducing phasiness. The amount of shuffling and the range of frequencies affected are user-controllable so that the effect may be optimized for a variety of stereo source material.

The S1 Stereo Imager also has the unique capability of rebalancing a stereo mix, and of altering the positions within the stereo stage of individual parts of a stereo mix. By allowing the user to manipulate the relative balance

of the central, left and right stereo components of a mix, it is possible to 'pan' the left/right and center components independently. In addition, the relative balance of the left/right and center signal may be also changed to either narrow or widen the perceived stereo image. S1's unique and informative stereo vector display allows the effect on level balance and positioning to be precisely visually monitored while making adjustments.

Stereo level metering is provided in both left/right and M/S modes so that the effect on stereoism can be fully monitored. As with some other Waves plug-ins, correction is provided for the two most common stereo imaging errors: channel polarity ('phase') errors and stereo channel reversal.

This stereo remix capability makes the S1 suitable not only for sweetening and enhancing adequately mixed stereo masters, but also for more drastic corrective and creative remix work. Because S1 doesn't rely on gimmicky processes to achieve its results, it may be used to make subtle improvements even to refined 'purist' stereo recordings. The S1 does not add any unwanted sound effects not already present in the original recording. Rather it is designed to enhance and rebalance what is already there.

S1 Component plug-ins

On most host applications and platforms you can select just the audio processing you need, and use only the power necessary to do the job. The S1 has several plug-in components in the menu. They are all *stereo-output-only* plug-ins and can not be used for mono-to-mono processing. On some native platforms, 88.2/96kHz support is now available.

Here are the current components as of this printing; new ones may be added in future updates, and will be noted in the plug-in Read Me files. Some examples for each component are listed:

- **S1 Shuffler** — This is the full plug-in, with all controls, including Width, Rotation, Asymmetry, Shuffling, plus all channel controls (phase, gain, etc.), and does not count clips; *(for mastering, to fix off-center or unbalanced mixes, precision image adjustment, change wideband width, expand low frequency image, Rotate true stereo sources without collapsing image, creation of strong antiphase content)*

- **S1 Shuffler (48)** — The full mastering plug-in (TDM only!), same as the S1-Shuffler, but 48-bit double precision resolution; *(applications same as S1 Shuffler but when increased resolution is needed)*

- **S1 Imager** — Omits the Shuffler and Bass trim controls. Ideal for adjust Width, Rotation, and Asymmetry only; *(saves processing power if Shuffling is not needed, excellent for stereo mixing to Rotate stereo-mic or submixes into place)*

- **S1 MS** — Simply an MS converter (a matrix) that converts left-right input into MS output, (or vice versa). It has no controls or metering, and is designed so that if the input does not clip, neither will the output. Due to the nature of the matrix, the gain is reduced by 3dB in this process to avoid clipping. Therefore, if you use two MS modules in a chain (say for equalization of just the M or S channel), then the total chain will have a nominal signal drop of 6dB (3dB per MS module, also assuming the EQ is flat with unity gain).

Chapter 2 - The S1 Interface

For complete information on basic Waves interface control, read the **WaveSystem Manual.**

The distinctive interface of the S1 lets you easily see and control the stereo soundstage of a file. The most noticeable feature of the interface is the **stereo vector display,** with its half-circle shape and soundstage graphic.

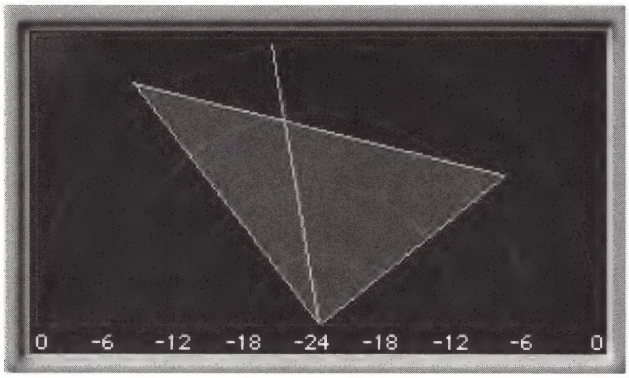

The stereo vector display of the S1.

The stereo vector display does not show anything about the input signal. It shows you:

1 - the output positions of the original left, center, and right sounds;
2 - an accurate indication of their changed levels in dB so you can adjust level-balance between original left, center, & right sounds.

The stereo vector display does not show the effect of **shuffling,** but only **gain, width, asymmetry, & rotation.** Shuffling is a process applied after these processes.

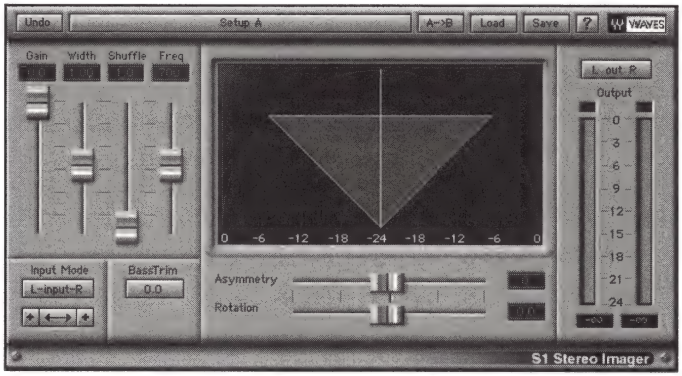

The S1 control interface.

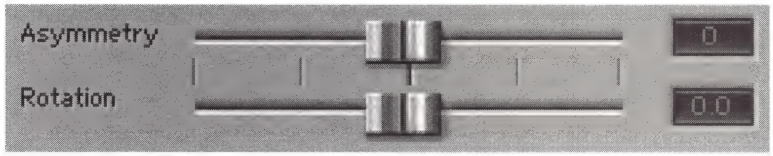

Gain - simply controls overall gain level without altering stereoism. Gain is displayed in dB.

Width - This alters the width of the input stereo signal. For width values less than 1, the sound stage is narrowed, becoming mono for width zero. For widths greater than 1, the sound stage is widened, with sounds at the two edges of the stereo stage being moved beyond the loudspeakers. The stereo is unchanged when Width equals 1. Start around 1.2 for moderate enhancement, plus a little Shuffling.

Asymmetry - This acts after initial width adjustment. This unique control does not affect central mono in-phase sounds in any way, but adjusts the relative level of left and right sounds. It differs from conventional balance control in that it keeps center sounds in the center. This control is particularly useful for altering the relative levels at the two sides of the stereo stage without moving a mix off-center, or biased in position to one side or the other.

Rotation control - This control allows adjustment of the centering of the entire stereo image without any effect on relative sound levels in different parts of the stereo stage. It is particularly useful for centering a stereo mix biased to one side of the stereo stage without altering the mix level balance, and for mixing stereo-miked sources, serving as a true stereo 'pan' control.

Useful tip: *No separate stereo balance control is provided, but if you drag both the Asymmetry and Rotation controls together, then their combined effect is that of a balance control. For most platforms, they may be dragged together by dragging in the space between the two controls, or by selecting both controls first then dragging on one.*

Shuffling - This increases stereo width at bass frequencies to help compensate for the fact that the ears hear stereo effects as being more narrow in the bass than in the treble. Subjectively, it has the effect of making stereo images more spacious. Unlike previous commercially available stereo shufflers, it is fully phase compensated, i.e. does not introduce unwanted phase errors between the stereo channels. This gives improved stereo sound quality with low listening fatigue. It is also carefully designed to minimize tonal alterations of the mix in the bass. It comprises two controls: shuffling and frequency. Shuffling controls the magnitude of bass width increase between 1 (no shuffling) and 3 (maximum bass width increase). The subjectively best value usually lies between 1.6 and 2.5.

Frequency - controls the frequency below which the Shuffling effect is increased. This frequency may be adjusted between 350 and 1400 Hz. For normal stereo monitoring, a frequency between 600 and 700 Hz usually sounds best, although a higher frequency may be better if used with the very close loudspeakers often used in multimedia, portable, or TV one-piece reproducer systems where the stereo speakers are in the same box. A lower frequency is useful when adjusting stereo originating in stereo microphone techniques where the microphones are spaced apart by perhaps 20 or 30 centimeters.

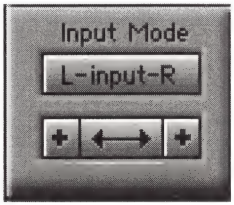

LR/MS input button - Changes input mode to accept either standard Left-Right inputs or Mid/Side (MS) inputs. See the chapter **About MS**.

Bass Trim - Adjusts the level of bass when the Shuffler is in use. When set to 0, bass level stays more or less the same. On some material, higher Shuffler settings may need bass trim of 1 or 2 dB boost to compensate for any subjective shift in bass level.

Reset - This button restores Width, Asymmetry, Rotation and Shuffling to settings at which they have no effect on input stereo. *Hint: By repeatedly using the* **Undo** *button after using* **Reset**, *you may do A/B comparisons between unprocessed and processed stereo effects, without altering gain or channel polarity and swap settings.*

Signal path (Order of action of controls)

The controls act in the following order between input and output.
LR/MS input button to set input mode;
then **Channel polarity** to correct possible stereo phase errors on the input signal;
then **channel swap** to correct possible stereo reversal on the input signal,
then **Width,**
Asymmetry,
Rotation,
Shuffling/frequency,
and finally **Gain.**

- First, preview or audition the stereo audio source you wish to process.
- Experiment with each control to learn the effects of each. It is best to return each control back to the default position before adjusting the next one! Refer to the section **The Controls** for detailed explanations of each.
- Due to the processing used in the S1 to modify the soundstage, you may sometimes need to reduce the gain level of any high-level files by 3 dB or more to avoid output clipping. For more information, see the **Output Clipping section.**

Graphical control

A unique and flexible feature of S1 is the ability to control the processing from the graph.

To control the S1 from the graph:

(For ease of explanation, the following instructions assume that you have clicked the Reset.)

- Click and drag on the center line (pointing straight up). This line represents the center channel.
- Drag left/right to change **Rotation.**
- Drag up/down to change **Gain.**
- Option-drag left/right to change **Width.** (PC uses Alt key).
- Option-drag up/down to change **Asymmetry.**(PC uses Alt key).
- *Hint for Mac: you can press and release the option key without re-clicking; this lets you change any of the four controls listed above in one step.*

Hints for Use

Using the S1 will provide subtle or extreme manipulation of the full stereo mix, stereo-miked tracks, reverb returns, submix groups, and so forth. The **Gain, Width, Asymmetry** and **Rotation** controls in combination can be used for a great variety of different adjustments and remixes of the original stereo source. (And the S1 Rotation will pan dual input mono, so can be used to replace the pan control across the mixing board for all channels, whether mono or stereo).

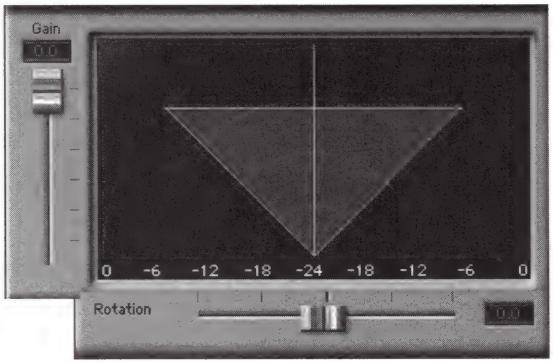

*Adjustment of **Gain** and **Rotation***

Quick tip: Use the "one-click" adjustment of an input stereo image by using direct graphic control, described on the first page of this chapter. You can hold and release the option(Alt) key while dragging without having to let up on the mouse to easily adjust all four controls; Gain/Rotation, or Width/Asymmetry.

To adjust stereo positioning without any effect on the internal level-balance of a mix, use the **Gain** and **Rotation** controls only (and the **channel swap** button if reversal of the i

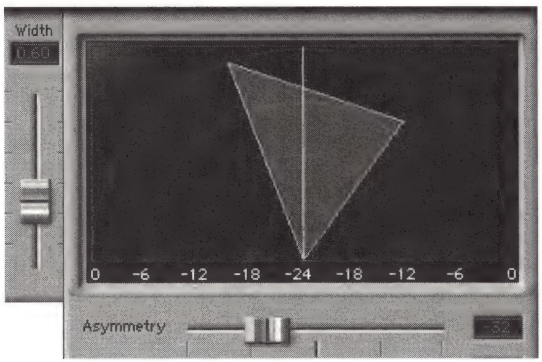

Adjustment of Width and Asymmetry to modify level balance.
Left, Center at 0 dB gain, Right at - 4 dB gain.

Conversely, to adjust the internal level-balance of a mix without any effect on the stereo positioning of central images, use the **Width** and **Asymmetry** controls only. This adjustment may be done by depressing the option (Alt) key and dragging the mouse from the vector display screen.

If the effect of a level-balance adjustment using **Width** and **Asymmetry** control is that the stereo image is too narrow, you may find that increasing the amount of **Shuffling** will help to restore width without much effect on perceived level balance.
Dragging **Rotation** and **Asymmetry** control together acts simply as a conventional *balance* control, albeit one acting after the **Width** control. (You can conceptually think of the Asymmetry control as being a balance control with a *rotation* to restore the center sound to center position.)

It is worth noting that the **Rotation** control has NO EFFECT on the internal sound balances of stereo mixes, only on their positioning.

Besides making the sound stage wider or narrower, **Width** does have an effect on level balance, making non-central sounds louder relative to central sounds as the width is made larger. This can sometimes be useful for changing the relative balance of central and edge sounds. But
beware: this only works in stereo, and does not change the balance heard by a mono listener, e.g. via a mono radio or TV.

Asymmetry changes relative balance of left & right both in stereo and in mono.

To mix multiple stereo sources using S1

Using faders with pan controls to mix several stereo submixes, or stereo-miked or other stereo sources can cause images to be narrowed or collapsed. This is unnecessary and awkward to adjust.

By simply using the **Rotation** control in the S1 plug-ins, a complete stereo input can be *rotated* into place. It's very similar to turning your head to listen off-angle, instead of using the pan controls, which has the side effect of reducing the stereo image. (This is done by panning the L and R unequally, such as L all the way to left and R being panned straight to center, 'squeezing' a sound halfway to the left.)

Insert the S1-Imager or S1-Shuffler onto each of your stereo submixes or sources you want to mix to a final output pair. Any mono instruments can be input too, and they will simply be panned, if input to both L and R of the S1 plug-in. Width and Shuffling have no effect on dual mono.

Mono Compatibility

When stereo sound is broadcast via radio or TV, a significant proportion of listeners will be listening to mono. In that application, it is important that stereo mixes also be mono-compatible, i.e. that no sound in the mix becomes colored, disappears, or becomes excessively quiet..

The S1-Stereo Imager is designed such that, with material mixed on standard stereo equipment, the processing does not cause coloration in mono listening.

The mono compatibility of reprocessed stereo material can be judged using the stereo vector display. The closer the "vector" or lines corresponding to left, center and right inputs are to being vertical on the display, the more mono-compatible is the mix, whereas the closer they are to being horizontal, the quieter will be that sound position in the mono mix. In practice, mono compatibility will usually be good if all three lines are within 45 degrees of being vertical, and adequate if all three lines are within 60 degrees of being vertical.

If mono compatibility is particularly important to you, you are advised to monitor the mix via a mono reduction of the mix as well as in stereo. Many mixing desks or monitoring systems have a "mono" monitoring button.

Chapter 4 - About MS and MS modes

What is MS?

In addition to conventional left-right input mode, the S1 also offers processing in what is known as MS stereo mode, also termed sum-and-difference mode.

MS mode is an alternative way of handling stereo signals that has many advantages and unique possibilities for stereo processing. Although MS mode has been used since the very start of stereo in the 1930s, it has only gained widespread use over the last decade, often being used in stereo TV production. But it's also used for various stereo microphone techniques and in other audio-only applications.

Briefly, M and S signals are derived from left and right by a matrix process. M stands for "mid" and S for "side" signals. The signal produced by an "MS" microphone is in this format and needs conversion from the MS to standard LR. However, the S1 offers unique processing of the signal before it is converted, and also allows any LR signal to be converted into MS for the same processing. In the **Special MS uses** chapter, examples of MS processing are explained.

The mid signal M, also termed mono, is simply obtained by taking the average of the left and right channels. Mathematically:

$$M = (1/2)(L+R)$$

where L is the left waveform magnitude and R is the right waveform magnitude.

The side signal S is obtained by subtracting the right signal from the left with the same overall gain as used for the M signal. Mathematically:

$$S = (1/2)(L-R).$$

The process used to derive the M and S signals from L and R is known as an *MS matrix.*

The MS matrix process may be applied again to the M and S signals to recover the original left and right, although an additional gain of 2 (+6 dB) is needed to get the original gain. Mathematically

$$L = M+S \quad \text{and} \quad R = M-S.$$

i.e. **Left** is the sum of the M and S signals and **Right** is the difference of the M and S signals.

How is MS used?

MS signals have various properties that make them useful for stereo processing. The M signal is simply the mono signal heard when a stereo signal is listened to on mono equipment. The S signal is zero when the left and right channel sound is identical as it is for sounds panned exactly to the middle of the stereo stage. This is why S is called the "side" signal —it only contains that part of the sound positioned to either side of the center position.

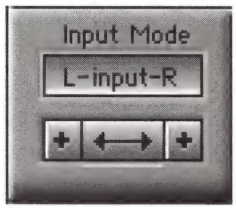

The standard S1 Shuffler or Imager can be used with input stereo signals from either standard left-right (LR) stereo or MS stereo. In both cases, the output is in standard left-right stereo mode. Click the button above the channel swap/polarity buttons to select which mode you need. In MS mode, what is normally the left input channel now accepts the M signal and what is normally the right input channel now accepts the S signal.

M and S outputs from an MS stereo microphone may be fed into the S1 using MS input mode, and the output will be in conventional left-right mode with the gain, width, asymmetry, rotation and shuffling controls affecting the sound in the usual way. MS signals may also be available from other equipment, such as stereo tracks recorded in MS mode on a video tape recorder.

Polarity/Channel Swap in MS mode

The polarity and channel swap buttons in the input to the S1 are intended to compensate for errors in the input signals. However, they work somewhat differently when the S1 is in **MS Input** mode! The polarity buttons alter the polarity of the incoming M or S signals as shown on the screen, and the channel swap button now interchanges the M and the S input signals.

The perceived effect of polarity and swap on the output stereo when the input is in MS mode is as follows:

Using only one of the polarity buttons in MS input mode has the effect of reversing the stereo image, swapping the left and right signal fed to the subsequent stages of S1 processing.

Using the channel swap button in MS input mode has the effect of leaving the left and right inputs to the S1 unaltered except that the polarity of one of them is inverted.

These effects arise from the mathematics of MS matrixing.

Therefore, be aware that if you wish to swap the left and right input while in MS input mode, you would activate the **M polarity button**, and not the channel swap button (as in illustration below).

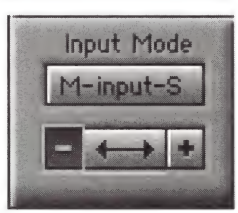

Use of MS meter mode

The S1 meters have two modes of display: LR and MS. The LR mode is standard Left and Right indication. MS mode is explained below.

To switch between LR and MS modes, click the L-R (or M-S button) above the meters. Switching the meter mode does not affect the audio output in any way.

MS meter mode shows the levels of the stereo "mid" and "side" signals M and S, i.e. the respective average of the two channels and the difference between the channels. The M signal contains most of the "mono" information in the stereo signal, and the S signal contains most of the stereo. Monitoring their levels give useful data about the stereo signal — the more level there is in S relative to M, the more stereo information there is.

The MS meter mode can be used to judge whether the input stereo material is in phase. If it is, the M level will generally be significantly higher than the S.

If the S level is generally higher than the M level, this indicates one of two possibilities:

(a) the most likely possibility, that the input channels are out of phase. For example, an analog-domain wiring error in balanced cables somewhere in the stereo signal chain, an incorrect setting of a phase switch, or possibly an incorrectly wired microphone capsule in a stereo pair of microphones. Correct this by altering one only of the two input polarity buttons to read " - " rather than " + ". Listen via loudspeakers to verify that this does indeed improve the sharpness of stereo images, particularly for sounds at the center of the stereo stage.

(b) in unusual cases (e.g. certain specialist one-point stereo microphone recording techniques such as MS or "Blumlein figure-of-eight" when presented with a very wide sound source stage, or if previous excessive width adjustment has been used) it is possible that the S meter can read higher than the M without a channel polarity error dues to excessive width. In this case, center stage sounds WILL be in phase, and channel phase adjustments should NOT be made.

This case is an unusual one - you may never encounter it. It can be distinguished from the case when the stereo channels are out of phase by one or both of the following listening tests. (A) narrow the width on the S1. Do center stage sounds become more solid and the sound become more natural? If so, then the original width was excessive. If, on the other hand, the center sounds become very quiet, the channels were out of phase. (B) If the rotation control is adjusted to the left, do center sounds also move over to the left? If so, then the channel polarity is correct. If the center sounds move over to the right, then the channel polarity is wrong.

If the M and S levels are roughly the same without processing, the two channels are largely uncorrelated and no judgement can be made from the meters about correct channel phasing. In this case, if there are no center stage sounds but the channels are entirely independent sounds, then no polarity adjustment is required. If there are center sounds, listening tests A or B above should be used to adjust channel polarity correctly.

When processing is applied to correctly phased input stereo material, it can be useful to monitor the results on the MS meters. In cases where an ambient effect is provided, it is often useful to adjust width and/or shuffling so that the S meter levels are comparable to or only a little less than the M meter levels. If the S meter levels are typically larger than the M meter levels, then the degree of width or shuffling or, just possibly, rotation is likely to be excessive, and should generally be reduced.

Chapter 5 - Other S1 controls

Some controls in the S1 are different or exclusive from those described in the WaveSystem Manual. These exceptions are covered in this chapter for convenience.

Clip meters and Level indicators

The peak and clipping indicators only give correct indications in L/R mode, where they show peak levels and any clips independently in the left and right channels. Do not rely on peak and clipping indications in the M/S mode. Whenever switching back to L/R mode, please reset the peak and clip indications to obtain an accurate reading.

Output Clipping (the Gain control)

The S1 can cause clipping of high-level files, even at default settings.

The unique phase compensation feature of the S1 assures low subjective "phasiness" of the stereo image. This uses phase manipulations to modify and improve stereo quality. However, even at the unity-gain default or reset settings of no shuffling or width alterations, these phase manipulations have a side-effect of altering the levels of peaks in the audio signal, so that on normalized or peak limited sound files, you may sometimes need to reduce the gain level by maybe 2 or 3 dB to avoid output clipping. Set the S1 gain so that no clipping registers on the output level meter in L and R mode.

For this reason, to get maximum output level, it is very strongly recommended to do normalization of level or peak limiting using the Waves L1 peak limiter AFTER processing by the S1 rather than beforehand.

Here is an example of how to use the S1 and L1 together for stereo image adjustment while maintaining unity gain.

Insert the **S1-Shuffler** first, followed by the **L1 limiter** or **+L1.** (This example is for systems that support multiple real-time processes, whether DSP or Native-based).

Assuming you have high-level files, or wish to be sure to avoid any clipping, **first set the S1 gain to -6.0dB**.

Next, set the L1 Threshold to -6.3dB, and the Out Ceil to -0.3dB. By using the L1 merely as a limiter to control any new peaks created by the S1 process, unity gain is maintained while giving 'safe' adjustment headroom for the S1 processes. Of course, additional gain beyond unity can be achieved with limiting, by lowering the L1 Threshold even further.

Chapter 6 - Special MS uses

Using the S1 MS options

The "Mastering" plug-in **S1-Shuffler (48) <for TDM only>** is preferred for high resolution mastering and critical applications. Its incorporation double-precision resolution ensures the highest possible quality in matrix operations and MS adjustments.

The smaller S1 component plug-ins are the preferred options where several Waves processors are cascaded one after the other or otherwise used together, because they then use up as little of the DSP power as possible, allowing more of other Waves component plug-ins to be used within the DSP limits of the user's system. They are ideal for MS use.

The power of the MS functions of S1 is greater when used with other Waves components, notably the Q10 (or Renaissance) equalizer, and various examples of use are described in the following sections.

MS Equalization

The small modules of the S1, along with those of the Q10, allow a kind of EQ that is different from conventional equalization of left and right stereo channels. It is known as *MS equalization.* This kind of EQ can achieve unique effects on stereo signals that cannot be obtained any other way, altering not only the overall tonal quality, but also the spatial effect of a stereo recording.

MS equalization is obtained by the following cascade of processors

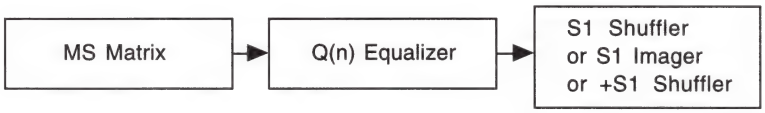

using an **MS matrix** followed by an **"n-band" stereo EQ** component from the Q10 (n= up to 10 bands of EQ), followed by an **S1 module** having the desired output facilities. If the output of the MS equalizer is to be fed to other processes in TDM or in 24 bits, the S1 Imager or S1 Shuffler modules may be used depending on whether shuffling is desired or not. One other option is to use another MS plug-in at the end instead of the Shuffler or Imager, but only if you don't want to use Width and other options in those components.

Setting up the MS EQ chain

When using any of these chains, first set the final S1 plug-in to MS input mode, and ensure that the polarities are both "+" and that channel swap is not activated. Reset the output S1 module to avoid any other S1 processing at first. In the Q(n) equalizer, set the polarity buttons to "+" and make sure the channel swap button (if present) is not activated.

The effect of this processing chain is that what is normally the left channel of the Qn equalizer now equalizes the M signal and what is normally the right channel of the Qn equalizer now equalizes the S signal.

By giving the M and S channel signals different gains at different frequencies, you can vary the width of the stereo in a frequency-dependent way. For instance, if you boost the S channel at 1kHz, then the output will have greater stereo width around 1kHz. If you boost the M channel at 1kHz, then the output will have a decreased width (and stronger 'phantom' channel) around 1kHz. A practical example would be to increase the midrange of Crown SASS microphone outputs, which many times have a 'hole-in-the-middle' effect on some sources; a reasonable boost around 400Hz of the M channel makes a significant difference without affecting the general stereo imaging in other frequencies.

You are now ready to start using the processing. Here are a few of the possible applications.

Uses of MS equalization

Overall EQ

Open the Q10 (or Renaissance EQ) interface. Ensure that the stereo channels are set to "Strap" or "Link". Then adjust the EQ in the conventional way to equalize both channels.

Frequency-dependent width

This is a new range of adjustments. Open the EQ interface. The procedure to adjust the width as a function of frequency involves the following steps.

[1] Start with identical EQ settings in both the M and S channels. Make sure the eq is in "Strap" mode.

[2] Select in a given band the type of filter you want, its frequency and Q, and any overall gain adjustment in that band that you wish. Then "unstrap" the EQ and click the **Right** button so that you can change just that side, which is the S channel.

[3] In that band now adjust the degree of boost and cut required to alter width in that band. Increase of gain will increase width and decrease of gain will decrease width, but only of that band's frequencies.

[4] In general, this adjustment will have altered tonal quality. To restore tonal quality, select "strap" and then adjust the gain in that band up and down until the tonal quality is what you want again. By adjusting both bands while strapped together, the new width adjustment is maintained (as a difference in gain between the two channels).

[5] If further adjustment is required, return to step 2.

It is important that you make these adjustment with strap on or off as indicated in each step. Otherwise the adjustments in the stereo sum and difference channels will no longer track together.

Please note it is preferable to do overall bass shuffling using the final S1 Shuffler rather than adjustments on the EQ, since the former is a design dedicated for the job.

Overall (wideband) width is also preferably altered on the output S1 plug-in, using the **Width** control. Altering the gain of the M and S channels of the Qn separately will also alter width, with an increase in S gain increasing width and a decrease decreasing width. But to avoid confusion, it is preferable to use the Qn gains only for overall gain adjustment and to avoid clipping in the Qn by altering both channel gains together.

One-point microphone compensation

For example, most stereo one-point microphones suffer from reduced width in bass due to low-frequency roll-offs in the directional microphone component characteristics. This may be compensated by boosting the S channel on the Qn at 6dB per octave by using a low shelf set to +18 dB and sliding it in frequency to compensate. Normally, the frequency will be set to below 100 Hz, a figure of 16 to 50 Hz being common for best results.

Similarly, dummy head recordings or *spaced identical* microphones pointing in the same direction (typical spacing 15 to 50 cm) can be "Blumlein shuffled" by the same technique but using a higher boost frequency, say + 3dB at 300 Hz on the S channel, along with say a -1 dB shelf on the M channel to compensate to the resultant exaggerated bass. The channels may then be relinked and the frequency slid up and down to adjust width or adjust for different microphone spacings.

FM broadcast noise reduction

FM broadcast reception typically has more noise in the S channel than the M channel. Material affected by S-channel hiss can be filtered in the S channel with a tunable notch filter to reduce noise, while having no effect on the tonal quality of center stage sounds, which are in the M channel.

One-sided Equalization

The small S1 modules with small Q10 modules can also be used for a unique mode of EQ termed one-sided EQ which allows the EQ on one side of the stereo stage to be adjusted without affecting the EQ of sounds in the middle of the stereo stage. For example, you can make sounds on the left have a brighter treble without affecting the sound of center-stage vocals. This can be very useful in remastering stereo mixes or submixes, as one has much finer control over the EQ effect in different parts of the stereo stage.

The basic process uses three small modules in cacade:

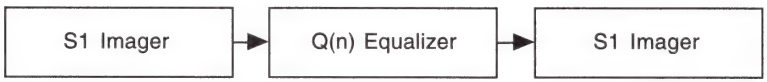

although the last stage can also be an S1 Shuffler or +S1 Shuffler if desired.

If one wishes to EQ the left side sound without EQ'ing the center, the first S1 Imager is loaded with the setup file "S1 LeftSideIn" and the second S1 Imager is loaded with the setup file "S1 LeftSideOut"

If one wishes to EQ the right side sound without EQ'ing the center, the first S1 Imager is loaded with the setup file "S1 RightSideIn" and the second S1 Imager is loaded with the setup file "S1 RightSideOut"

The input and output S1 imager files are designed such that if the Qn is set to flat, then apart from an overall gain change, there is no effect on the stereo. If the channels of the Qn are strapped and the EQ's varied together, this adjusts overall EQ in the usual way.

What is new here is what happens if one unstraps the two channels of the Qn. When using the LeftSide setups, adjusting the EQ of the left channel of the Qn adjusts the EQ of the left-side sounds of the input stereo signal, but has no effect at all on sound at center stage, which remain unequalized and stay at the center position!

When using the RightSide setups, adjusting the EQ of the right channel of the Qn adjusts the EQ of the right-side sounds of the input stereo signal, but has no effect at all on sound at center stage, which remain unequalized and stay at the center position! Excellent when needed.

Chapter 7 - More about shuffling

Shuffling is a stereo enhancement process that has few unwanted side effects, but which can add a sense of spaciousness and life to an otherwise ordinary stereo effect.

Shuffling is particularly suited to recordings made using standard stereo panning techniques made on most stereo mixing consoles or their digital equivalents.

The stereo effect from two loudspeakers creates the illusion in the ears and mind of the listener of illusory phantom images lying between the loudspeakers. This illusion, however, is not perfect, and one known fault is that for sounds panned to intermediate stereo positions, the apparent stereo stage width is narrower at bass frequencies than in the mid and high treble. This is known to be connected with the fact that the ears and brain determine sound direction by different methods below and above around 700 Hz.

Shuffling is a process that widens the bass frequencies so as to make the phantom bass and treble images seem to be of comparable stage width. This generally occurs when the shuffling level is set at around 1.6 and shuffle frequency around 650 Hz. This setting is recommended for sharpest stereo imaging.

If you increase the bass width even further, so that the bass is over-wide compared to the treble, it will create less precise stereo images, but generally tends to create an illusion of "spaciousness" and sounds from beyond the two stereo loudspeakers. This helps makes up for the fact that the stereo image is normally confined to the space between the loudspeakers, rather than being all around as in real life. This effect is generally best when the shuffling level is between 2.0 and 2.5, and shuffle frequency around 650 Hz. This setting is recommended for best spaciousness on most material.

Such settings are especially useful on stereo material that otherwise sounds too narrow.

Shuffling has no effect on the stereo imaging of central or mono images, but only on sounds panned away from the center. Its effect is not to enhance mono images, but rather to enhance the stereo effect already present in stereo images.

Unlike previous commercially available stereo shufflers, it is fully phase compensated, i.e. does not introduce unwanted phase errors between the stereo channels. This gives improved stereo sound quality with low listening fatigue. It is also carefully designed to minimize tonal alterations of the mix in the bass. The application of shuffling has no significant effect on mono compatibility.

It comprises two controls: "shuffling" and "frequency".

"Shuffling" controls the magnitude of bass width increase between 1 (no shuffling) and 3 (maximum bass width increase). The subjectively best value usually lies between 1.6 and 2.5.

"Frequency" controls the frequency below which the shuffling width is increased. This frequency may be adjusted between 350 and 1400 Hz. For normal stereo monitoring, a frequency between 600 and 700 Hz usually sounds best. A higher frequency may be better if the sound is intended to be used with the very close loudspeakers often used in multimedia or portable or TV one-piece reproducer systems where the stereo speakers are in the same box. A lower frequency is useful when adjusting stereo originating in stereo microphone techniques where the microphones are spaced apart by perhaps 20 or 30 centimeters.

Chapter 8 - History of stereo controls

Width control was invented in 1931 by Alan Blumlein. It is probably the oldest control for adjusting stereoism, pre-dating even the balance control. It has been included in many stereo mixing desks, and even a few hi-fi amplifiers, since then. However, despite its usefulness, it still has not come into universal use.

Asymmetry control: this is the first commercial product using this control, which was devised by Michael Gerzon in the mid 1980's.

Rotation control: Although a stereo rotation control was available in the 1950's from at least one central European manufacturer, Telefunken, it never came into general use and has been forgotten except in German broadcast mixing desks. An Ambisonic version of this control has been available in Soundfield microphone control units manufactured by Calrec, AMS, Drawmer and AGM since the late 1970's.

Shuffling. The terms "shuffler" and "shuffling" were coined by Blumlein in the 1930s. This particular kind of shuffling was devised at EMI in England by Vanderlyn in the 1950s, and used for a brief period by EMI in their early "Stereosonic" recordings, but fell out of use, largely because at that time EMI used recording techniques for which it was not especially effective. Individual recording engineers, including Tony Faulkner in England, devising their own equipment since then have used shuffling, and it was rediscovered by David Griesinger at Lexicon in the mid 1980's and incorporated as an effects program in various of their digital effects units. Although the importance of phase compensation was known to Vanderlyn in the 1950's, and was incorporated into EMI's own equipment of the period, the S1 is believed to be the first commercially available shuffler to incorporate phase compensation.

Further reading on stereo

Probably the best generally available source of information on the enormous amount known about the stereo effect can be found in:

"Stereophonic Techniques", edited by John Eargle. An anthology of reprinted articles on stereophonic techniques, Audio Engineering Society, New York, 1986. Obtainable from the Audio Engineering Society. This anthology contains many of the classic articles and research papers on the stereo effect from 1931 to 1985.

Chapter 9 - What kind of stereo will the S1 work on?

The S1 will give excellent results for adjusting the stereo effect of most kinds of standard stereo material produced in the recording studio and on conventional stereo mixing desks or their digital equivalents.

For such stereo, the only precaution needed is to ensure that the two channels are in phase. Accidental polarity reversal between the stereo channels can sometimes occur due to miswired leads or connectors in the analog domain. If stereo channels are out of phase, the S1 will not work correctly until the phase is corrected by the polarity buttons (see the section **phase reversal** and **channel swap** earlier in this manual). You may notice a polarity reversal because central stereo images are "unsharp" and bass-light in sound, or by the fact that when the output meters are in M/S mode, the level on the difference S channel is consistently larger than the M output level. With such out-of-phase stereo, you will find that center-stage sounds move in the wrong direction as you apply the rotation control. To correct this problem, click on one of the polarity buttons so that it reads "-" while the other remain at "+".

However, for some stereo microphone techniques used for live recording, the S1 may not always work in quite the way expected. The basic advice here is simply to use your ears in judging the S1's effect when used with such recordings. The S1 will usually work very well with recordings made with "one-point" stereo microphones in a single small housing, sound field microphones and with MS stereo microphones.

However, when the microphones in a stereo pair are spaced apart by more than a few centimeters, sounds that are not exactly center-stage will arrive at different times at the 2 microphones, and in different phases. These time and phase difference will affect the operation of the S1.

When used with stereo microphone systems where the mics are spaced apart by more than a few centimeters but less than about 1 meter, you may find that the width control does not always work well on the "phantom" sound images between the loudspeakers, and can sometimes gives such images a colored and more out-of-focus sound. In such cases, leaving the width near a value of 1 and using only the gain, asymmetry and rotation controls may work best. You may also find it helpful to use a lower value of shuffler frequency than normal in such cases.

With stereo recordings using very widely spaced microphones (say more than 1 meter apart) as the stereo pair, you may find that rotation or asymmetry control used on their own are not effective, but that dragging them together, so that they act as a balance control, may be effective. In this case, width adjustment may again sometimes prove effective.

The S1 may not prove very effective on recordings made with a "binaural" dummy head, or similar approaches to recording such as the Schoeps "sphere" microphone or the Crown SASS system. These systems produce a very non-standard kind of stereo, and may not be improved by the S1. These systems ideally require the use of a different kind of shuffler to that incorporated into the S1 (further details on this may be found in M.A. Gerzon, "Applications of Blumlein Shuffling to Stereo Microphone Techniques", J. Audio Eng. Soc., vol. 42 no. 6 pp. 435-453 (June 1994).). Width control may not be effective for such stereo microphones. However, although the S1 was NOT designed for use with the stereo produced by such microphone, if you like an effect it produces on them, feel free to use it.

In the Stereo Vector Display:

- **(no key)** to control Rotation (X-axis)
  and to control Gain (Y-axis).
- **option-drag** to control Asymmetry (X-axis)
  and control Width (Y-Axis). (PC uses Alt key).

- For complete information on interface controls, read the **WaveSystem Manual**. (PCs use the Alt key instead of Option, for example).

Chapter 11 - Stereo monitoring

To monitor stereo effects correctly, it is important to make sure that your stereo monitoring is set up correctly - in particular that the stereo monitor speakers are matched and are in phase, and that the left and right channels are the correct way round!

A "left/right" sound file, or the channel identification track on a test CD, can be used for getting them the right way round. Any speech or music mono sound file, especially those with a lot of bass, can be used to get the speaker polarities correct — the image will be much sharper and with more bass if the speaker polarities are correct.

If stereo phase is wrong, you may correct it by reversing the "+" and "–" wires from the stereo power amplifier to the loudspeaker ON ONE STEREO CHANNEL ONLY.

If your stereo amplification system is fed via balanced cables and your speakers are out of phase, it may be that one of these cables is not wired correctly, causing a monitoring stereo phase error. If so, the connecting wire should be replaced or rewired to ensure the same polarity.

If the loudspeaker is bi-amplified, i.e. different amplifiers are used for different frequency ranges, it is important that ALL the amplifiers on that channel are used in the same polarity. If in doubt about this, consult the manufacturer of the bi-amplified loudspeaker or a specialist dealer for expert advice on this.

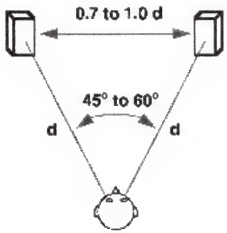

The best stereo monitoring will be obtained if the loudspeakers are set up symmetrically towards the left and right, with the listener placed symmetrically so that he or she is equally distant from them. Normally, stereo monitoring is optimal if the angle subtended by the two loudspeakers at the listener is around 45 to 60 degrees, i.e. if the distance between the loudspeakers is between 0.7 and 1 times the distance of the listener from each of the loudspeakers.

The subjective stereo effect heard via headphones will generally not be the same as via loudspeakers. If it is important to know the effect heard on both loudspeakers and on headphones, both should be used to check a stereo mix effect. One cannot with certainty predict the effect for headphone listeners by loudspeaker monitoring or vice-versa. In general, reverbs, and other uncorrelated materials will sound louder in headphones, sometimes by 6dB or so.

Q10
Paragraphic Equalizer

Table of Contents

Chapter 1 ..... About Q10 2

Q10 component plug-ins 2
Factory Presets 2
Setup library 3

Chapter 2 ..... Using Q10 4

Color graphing 4
The Equalizer 5
Q Control 6
About Q 6
Freq 6
Gain 6
Dual axis control 7
Filter type 7
In/out button 7
How to use the EQ controls 8
Direct graphic control 8
More control options 9
In the graph, on any Band Marker(s): 10
Arrow Keys 10
Clicking on filter type buttons: 10

Chapter 3 ..... Internal clipping .. 11

Chapter 1 - About Q10

Our challenge in designing Q10 was to develop an equalizer that gave excellent sonic performance and ease of use, but we also wanted to build in enhancements and facilities that would be impractical, if not impossible, to implement in the analog realm. The Q10 also provides a visual representation of the EQ curve that you can adjust by 'dragging' the curve itself so you don't have to rely on conventional controls. However, if you know what EQ settings you want, you can simply enter the numerical values to achieve the same result.

A traditional parametric equalizer might incorporate three or four bands of equalization, Q10 provides up to 10 bands of equalization per channel, configurable as either two independent mono equalizers or as a precisely ganged stereo equalizer. Use just what you need for the job. Each of the 10 bands can be true parametric, high/low shelves, or high/low pass (cut) filters, and to ensure optimum fidelity, each is noise-shaped for the best possible signal-to-noise ratio.

New 88.2/96khz support for native

If you are using Q10 in a native application which supports the higher sampling rates, the Q10 now supports these sampling rates, although the frequency scale does not change to show these frequencies. Therefore, if you are using a shelf or bell that would normally have "dropped" near the window edge, now the filter shape will continue onward toward the now-higher Nyquist!

Q10 component plug-ins

You can select just the audio processing you need and use only the power necessary to do the job. Use only as many bands of EQ as you need for efficient use of your platform's processing capabilities. The WaveShell makes the most of your system.

Depending on your platform, the Q10 has several plug-in components in the menu. You can select up to 10 mono or stereo EQ bands ; each plug-in is appropriately named for the number of bands, i.e., Q4, Q10. When being used in Native applications, if you don't have a band turned on, the CPU is not being used. In the case of DSP systems (such as TDM), then you should select the number of bands that best fits your needs in order to conserve processing power.

Factory Presets

These presets will show up in all components, and each requires different numbers of bands to work. The Preset name has a small number in brackets (<4>) at the end to show how many bands are needed.

Example:

if you see **SuperNotch (-46dB) <4>** in the menu,

it says you need at least 4 bands to use this 46dB deep notch filter.

Setup library

A wide selection of configurations are available in the Setup Library providing instant access to complex filters, which may be used as they are or modified for specific applications. Categories of Setups include Multimedia, PseudoStereo, Band-limiting, Pre/De-emphasis, Plateau filters, Tilt filters, Crossovers, Harmonic combs (hum removal), and so on. The Library includes comprehensive printed documentation which offers practical hints and insights into the use of each setup in the Library, as well as historical and practical use of equalization. We suggest that power users read it and enjoy! It was written by the late Michael Gerzon and has many amazing and powerful hints. We didn't edit or revise the Setup Guide document, so some of the examples and materials may refer to Sound Designer II, the editor Michael used to create the setups.

Most of the Setups are user-adjustable and applications include corrective applications, such as hum or hiss reduction, as well as creative effects such as AM radio or telephone simulation, pseudo-stereo, comb filtering and so on. There are also several utility filter configurations including brickwall crossover filters and other special equalizer characteristics.

Chapter 2 - Using Q10

This chapter gives a quick overview of how to use the Q10.

For complete information on using the interface of Q10 and any Waves plug-in, read the **WaveSystem Manual**. Some of the aspects of these controls *unique to Q10* are discussed in this chapter.

In brief, parameters may be entered either from the keyboard or by using the mouse as a slider, and direct control over the EQ curve is possible by using the mouse to 'drag' points on the curve directly.

Color graphing

Q10's display comprises a single window which makes use of color to present the maximum amount of information in an easy to understand manner. Because both left and right audio channels are displayed in a single window, the left and right equalization curves are represented by blue and red lines respectively, and for those few using mono monitors in 1-bit mode, the left channel is represented by the darker gray line, while the right is represented by the lighter gray line.

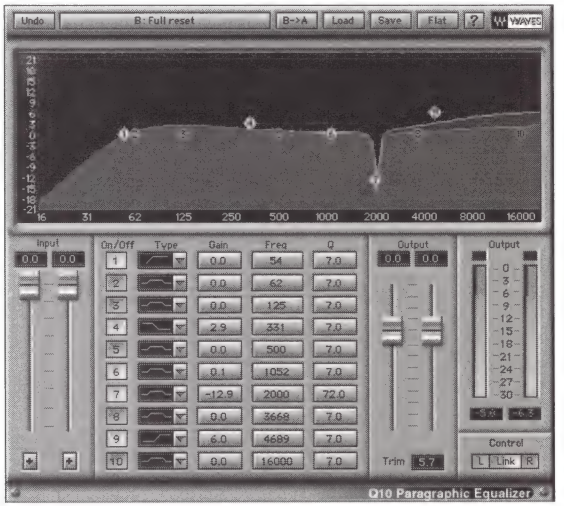

The Equalizer

Each of the 10 bands of EQ provided by Q10 has exactly the same controls and the same parameter range. The only difference is in the default frequency values.

The parameter values indicated in the small Value Windows may be changed either via the

keyboard or using the mouse. See the **WaveSystem Manual** for general details about using the controls. Exceptions specific to Q10 are discussed later in this chapter.

Value Windows of the Q10

Q Control

The area beneath the frequency curve window is a table with Value Windows for Frequency, Gain and Q for each of the 10 equalizer bands, along with the band's filter type and In/Out button. The right hand column Value Windows show the Q value which corresponds to the width of the frequency range for that band. The higher the Q setting, the narrow the band of frequencies that is affected.

About Q

Q is a way of expressing the frequency width of a filter in relation to the center frequency of the filter band and may be represented by the mathematical relationship: Fc/Fw (center frequency divided by frequency width at the traditional '-3dB point', both in Hertz). For example, a Q of 2.0 at 1,000Hz gives a bandwidth of 500Hz (1000Hz divided by 2).

In the Q10, the filter design adopts a constant bandwidth rather than constant Q approach which means that the literal interpretation of Q is only strictly applicable at maximum boost setting. At other gain settings, the

width of the 'bell' is kept nominally the same as at maximum boost, not because we want to be awkward, but because we feel that this approach makes the control of the equalizer more intuitive and more musical. In the Q10, when using small boosts and cuts, the Q value shown is approximately 5.5 times greater than the traditional engineering definition of Q. However, the basic principle of Q still applies; the higher the value of Q, the sharper the filter and the narrower the range of frequencies that is affected.

For typical EQ applications, Q values of between 0.5 and 6 are typical, but for specialist applications, for example, notching out specific frequencies (such as hum), much higher values of Q (up to 100) may be used. Figure 2 shows examples of both high and low Q settings. Note that because Q10 uses a variable Q filter topography, for most small boost/cut uses, these typical 'engineering' Q values of 0.5 to 6 would be shown as approximately 2.7 to 33 Q in the Q10.

Freq

The frequency control (Freq) Value Window displays the center frequency for each band when bell or bandpass mode is selected. When in low/high pass or low/high shelf modes, the Freq parameter controls the corner frequency of the filter. Traditionally, the corner frequency of a filter is defined by its -3dB point, but in the Q10, this is only strictly true at the maximum boost setting. At other gain settings, the constant bandwidth filter design maintains a nominal Q which is roughly the same as at maximum boost.
The Frequency control values are adjustable in one-sixteenth-tone steps (96 steps per octave), rounded to the closest integer. The left/right arrow keys may be used to step the selected Freq value in one-sixteenth-tone increments.

Gain

Gain may be adjusted in the range -18.0 to +18.0 dB, in increments of 0.1 dB via the Gain Value Windows or by 'dragging' the EQ curve. The displayed frequency response graph shows a range of +/- 21 dB.

Dual axis control

When setting up an equalizer, many engineers will adjust the Q or frequency with one hand while varying the gain with the other. Q10 goes one better by allowing the user to employ both the Y and X movements of the mouse to access two sets of parameters simultaneously, using just one hand. Selecting either the Frequency or Q of any band links that parameter to horizontal movements of the mouse while Gain is changed by vertical movements. Any change in parameter value will be reflected in the shape of the resulting EQ curve, and similarly, any adjustments made by dragging points on the curve will cause the Value Windows to change accordingly.

When adjusting the EQ curve directly, the X axis controls Frequency (Freq) by default while the vertical or Y-axis movement of the mouse adjusts the gain. Holding down the keyboard Option button at any time switches control from frequency to Q (PCs use the Alt key). Option-clicking on the band marker will not reset the values; option-clicking on a *value window* such as Gain or Freq will reset it to its default value.

Filter type

The FilterType Value Window, selects the type, or mode, of filter used for each band. The selected Type is indicated by means of an icon representing the filter curve.

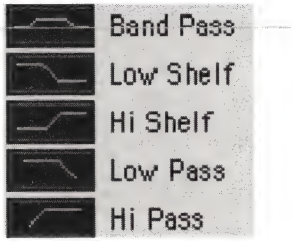

The default Type is the Bell or bandpass filter which is used to provide parametric equalization or to construct notch filters. Clicking on the Type icon will step through the filters in the following order: Parametric (Bell), Low shelf, High shelf, Low pass, High pass. You can also option-click on the filter type button to see a popup menu for direct selection.

In/out button

Band In/Out buttons

To the left of the Type button is the In/Out (on/off) button for each band. Clicking on the In/Out button causes the button to 'light up' and the filter band number to the left of each button turns red when the corresponding band is selected for adjustment. This makes it easy to determine which band (or group of bands) has been selected, especially when the band markers are in close proximity to each other, as is often the case when using the Setup Library or tweaking a complex filter.

If a point on the filter curve is selected with the mouse, the band number beside the In/Out button turns red. This makes it easy to determine which band is being adjusted in complex setups. The number turns red whether the band is switched In or Out. Note that when a filter band is bypassed, moving the filter's 'handle' on the frequency display will have no effect on the shape of the frequency curve until the filter is switched back in.

How to use the EQ controls

Perhaps the most intuitive way to adjust the equalizer is simply to use the mouse to modify the EQ curve display. Not only does Q10 allow you to control single bands using the mouse, it is also possible to select multiple bands for simultaneous adjustment when necessary. You may take control of a band with the mouse and use the two axes of movement to change either the Freq and Gain or the Q and Gain simultaneously. You may also bypass single bands or selected groups of bands. Please read the WaveSystem Manual for complete details of the controls.

The following example will help you to get the feel of setting up the filters by using the Value Windows and the mouse. Select a suitable section of audio, then click on the Preview button if appropriate for the host application, or pass audio through the system if the host application permits real-time control. If you are using Preview mode, you should now hear your short selection of audio looping continuously. The maximum length of the loop is usually determined by the amount of computer RAM available to the host application in which Q10 is running. Where real-time processing is supported, you will be able to adjust Q10 as the audio plays through normally. See your host application manual for details.

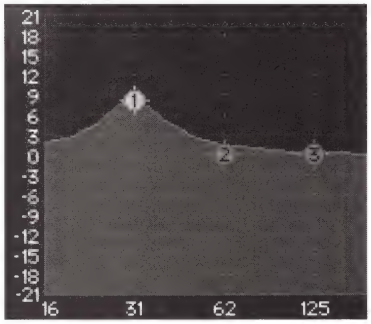

Band 10, showing +8.0dB Gain, controlled from Value Window

- First, click the In button of Band 10, the highest frequency band; the In button is located at the lower edge of the screen. It will turn red, indicating that this band is now switched in. Note also that a Band Marker has appeared on the flat line of the frequency curve graph .
- Click and hold on the Gain value for this band, and slowly move the mouse forward ('up' in the y-axis) until you have introduced around 8dB of gain. You should hear now hear the result of this high frequency boost, and you will notice that the Band Marker moves as you change the parameter value. When you release the mouse button, the graph will be redrawn and you'll see the EQ curve change shape to reflect the actual response curve being created by your settings. The Band Marker is at the top-center of the bell curve you have created and represents the center frequency (Freq) of this band.

- Now click and hold the Q control for the same band, and slowly move the mouse to the left (to control Q, you move the mouse horizontally) so that the Q value gets larger and the filter response becomes narrower. Set the Q to about 10.0. Again notice that the shape of the filter bell curve changes once you release the mouse button.

- Click and hold on the Freq control for this band, and slowly move the mouse back and forth, from left to right, and you'll hear the center frequency of the filter change. (Remember, to control Freq, move the mouse horizontally).

Direct graphic control

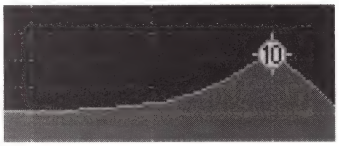

- Click and hold on the Band 10's Marker in the equalization graph and drag it up and down using up/down mouse movements to change the Gain value. Without releasing the mouse button, move the mouse from left to right and notice that the Freq value changes accordingly. If you like, you can move the mouse diagonally and change both Frequency and Gain values simultaneously.

- To change Q with horizontal mouse movement, hold the keyboard Option(Alt) key down while moving the mouse. Vertical movement will still change the Gain of the selected band, but horizontal movement will now control Q. Note that when modifying parameters, the mouse movements are the same in the Frequency graph as they are in the Value Windows, and because the Band Marker moves in the same direction as the mouse, you can think of the action simply as dragging the Band Marker.

- Multiple bands may be selected and manipulated simultaneously by holding down the Shift key as you select the Band Markers you wish to control, or by drag-selecting a group.

- Bands may be 'toggled' In and Out by a clicking on a Band Marker while holding down the Command key (PCs just doubleclick the band marker).

Being able to control the values of a band directly from the graph is a powerful advantage as it allows you to set up the EQ by feel rather than by numbers. Bear in mind that small EQ changes can make a significant difference to the way something sounds, so don't feel obliged to make use of the Q10's full range unless you really need it!

More control options

- Change the filter type by using an Option-Command-click on the Band Marker (PC doesn't support this).

- 'Lock' into one parameter without changing the other by holding the Control key down while moving the mouse. For instance, either before or after clicking on a Band Marker, press the Control key down. Then, move the mouse to adjust Freq without changing the Gain. Release the Control key. Then, press it again, and this time move the mouse to adjust the Gain without changing the Freq.

Note that direct keyboard control can only be used if supported by the host application or platform. Some platforms will not support all keyboard commands, and most digital sequencers don't allow use of the keyboard at all by plug-ins. Check your host application manual, plus any Read Me files about your host that came with your Waves package.

For platforms that do support keyboard commands, here is a summary of the Q10 commands (some are unique to Q10, all others are covered in the **WaveSystem Manual**).

In the graph, on any Band Marker(s):

option (alt)-click	to control Freq and Q.
click	to select a band (or any Value Window), and *drag* to control Freq and Gain.

Note: **You can press the Option key at any time to control Q without having to click again. Release the Option key and control returns to Gain.**

control-click	to adjust only one parameter. (The first move becomes the active one.)

Note: **You can press the control key at any time to activate the parameter lockout. Then, release to allow changing of both.**

command-click	to toggle band in/out.
shift-click	to select or de-select a band
option-command	to toggle filter type

Arrow Keys

To adjust any selected value(s) by small increments, use the arrow keys. (Note: The increment for each value is smaller than what will be displayed, so it may take several keypresses to see the value change on-screen.)

Left/right arrows control Freq or Q, and up/down arrows control gain. Pressing an arrow key once will change the selected value(s) by the smallest possible increment. Press and hold for continuous change.

Clicking on filter type buttons:

option(alt)-click	to see popup menu of filter types for direct selection.
click	to toggle through the filter types

Chapter 3 - Internal clipping

To copy/paste selected buttons, grab markers, etc.

Select the controls you want to copy (shift-click, drag-select, etc.)
Press the letter 'c' to copy the selected value(s).
Press the letter 'v' to paste the selected value(s).
Undo is active for any accidental paste action.

If you have a band selected in the graph (a Band Marker) and you press 'c', all three parameters (Gain, Freq, and Q) are copied. All three can be pasted to another band, even in the alternate setup (for example, copied from a band in Setup A and pasted into a band in Setup B). Multiple bands can be copied and pasted simultaneously. They are copied in order and do not have to be contiguous. For example, you could also copy from bands 1, 2, and 3, then paste to bands 4, 7, and 9, respectively.

Note: **It is possible to copy/paste from any control to any other control, which may yield unexpected results (such as pasting a Q value into an input fader). If you try this, be aware that what is actually pasted is an internal value and not the visible numerical value. The internal value is scaled for each control's range.**

Under certain circumstances, internal clipping can occur. With high level files (already normalized, or with high peaks), several situations can occur that may cause peaks that did not exist in the unequalized signal. This is true even for filters that have no gain, such as a high-pass filter, or for filters with a cut. The VU meters are output meters and will not show this internal clipping. To show this kind of clipping, meters would be needed between each band of the equalizer. Therefore, if you are using a soundfile with high levels (either normalized or very near 0 dbFS), just reduce the Input Faders 3 to 6 dB to allow for internal peaks.

Q10
Setup Guide

Table of Contents

Introduction 2

Chapter 1 ..... Pseudostereo 6

Chapter 2 ..... Crossovers .. 19

Chapter 3 ..... Bandlimiting .. 33

Chapter 4 ..... EQ Tools .. 40

Chapter 5 ..... Harmonic combs .. 46

Chapter 6 ..... Multimedia .. 50

Chapter 7 ..... Distortion effects .. 55

Chapter 8 ..... Tilt and Semi-Tilt .. 59

Chapter 9 ..... Pre-/De-Emphasis .. 64

Chapter 10 ..... Plateau .. 72

Chapter 11 ..... Stereo Guitar .. 76

Chapter 12 ..... Left/Right Setups .. 77

Chapter 13 ..... Bass separation .. 78

Introduction

Because of Q10's 10-section filter topography and its facility for simultaneous control of several EQ bands or parameters, it may be used to create special effects and to produce spectral manipulations that would either be extremely difficult or impossible using conventional equalizers. Q10's 'select and drag' control interface makes the control of even very specialized EQ settings very straightforward, but specialized EQ settings can be very time consuming to create. A complete Q10 stereo filter setup involves the fine-tuning of 60 filter parameters, and in a busy commercial environment, there is always the temptation to manage without a certain effect, simply because there's no time to spend setting it up. That's where the Waves Q10 Setups Library can help you.

This Library contains a great many filter setups, optimized for specific applications, and these may simply be loaded in and used just as they are, or fine tuned to the job at hand. Many of the settings make use of Q10's multiple selection facility allowing a complete filter setup to be adjusted in one click or drag.

This manual describes the various filter setups in the Library and provides information on how they might be used, and on how they may be adjusted or modified. This manual has been termed a 'Tutorial Manual' because it does more than furnish minimal instructions; it helps users unlock the power and potential of the Q10 to achieve the ultimate in creative and corrective EQ effects.

Using these setups, the Q10 is shown to be more than just a conventional parametric equalizer, it becomes a unique effects unit and processor.

All effects (with a few noted exceptions designed for use with sample rate conversion) will work at both 44.1 and 48 kHz sampling rates.

Please let us know how this library serves you, and of other setups you'd like to have included in future releases.

Operational aspects

Clipping distortion & Overload

The input and output gains of setups are generally adjusted so that the overall energy gain is unity to permit direct A/B comparisons between sounds in effected and bypass modes in preview. However, be warned that with audio files running to near peak levels, this can sometimes result in clipping or overloading at the output, causing audible distortion. Any overloads will be registered by the overload indicators above the level meters. Reducing the input levels by around 3 dB will reduce the risk of overloading.

If overload has occurred on a processed audio file, you should undo the processing, reduce the input gain, and process again.

The input gains in most setups are generally adjusted so that for signals with a normal frequency distribution, the internal processing of the EQ will be free of internal clipping distortion. Most of the time, overload distortion will occur only at the output and can be avoided by pulling down output gain. However, to be absolutely safe when processing signals with an abnormal frequency content, it is better to perform any needed gain reduction using the input faders.

An exception to this rule is the intentional distortion effects!

Sampling rates

Except where specifically stated otherwise, the effects here were devised primarily for use with audio files sampled at 44.1 or 48 kHz. Many of the setups will not work in the intended fashion at other sampling rates, especially at significantly lower sampling rates such as the 11 or 22 kHz multimedia sampling rates.

(In the older Q10 version 1.0, the setups will not load correctly at sampling rates below 44.1 kHz, and the graphical display will be incorrect.)

The setups have been optimized for use with files sampled at 44.1kHz. If files of a lower sample rate are processed, the filter frequencies will be scaled down proportionally to the sampling frequency. For example, all filter frequencies will be halved when the setup is loaded to process a 22.05 kHz sampling rate file; the new frequencies will be displayed correctly in the frequency value windows and on the graphical display.

Effects that will work at all sampling rates include EQtools, Supernotch, Superparametric EQ, Harmonic Combs, and Pseudostereo and Crossover effects, although their actual frequencies will be changed. All but some pseudostereo effects may be altered in frequency to compensate for the difference in sample rate as described below.

The bandlimiting and distortion effects will generally come out at the wrong frequencies at lower sampling rates, but some (including **Q10 telephone**) may be re-adjusted to restore the correct frequency range.

Important control information

Select-and-drag control

The Q10 permits a large number of equalizer parameters to be 'linked', i.e. controlled at the same time with a single movement of the mouse. See the main Q10 manual for more information on this facility.

For most adjustments on stereo material or on mono inputs, the stereo channels should be linked.

Example

Load the Q10 EQtools, and click the 'Open' button. The screen shown below should appear.

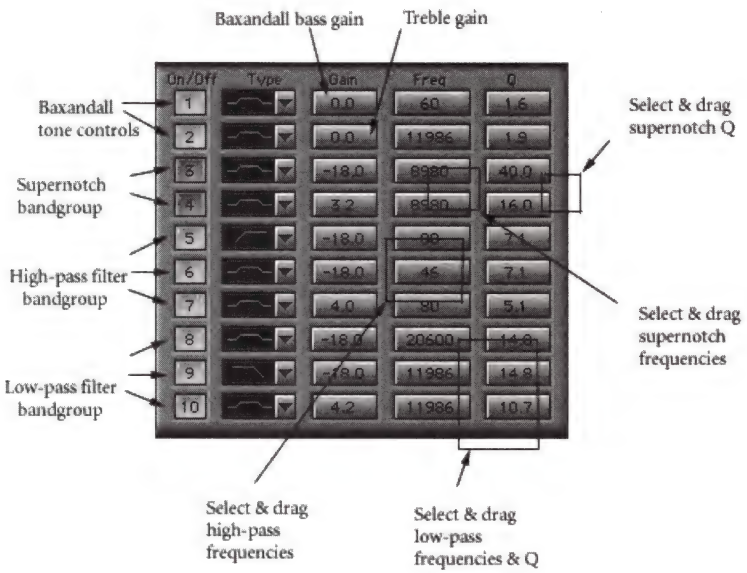

The select and drag options and controls of the setup Q10 EQtools.

Note that many of the setups use several Q10 bands to achieve a single filtering effect.

For example, the **Q10 EQtools** setup file when loaded provides four different equalization tools. Bands 1 and 2 provide the respective bass and treble of a Baxandall equalizer while Bands 3 and 4 provide a 'Supernotch' filter for removing unwanted narrow frequency bands. Bands 5, 6 and 7 provide an adjustable steep-cut, high-pass filter for removing lowbass frequencies. Bands 8, 9 and 10 provide an adjustable steep-cut low-pass filter for removing high treble frequencies. Each bandgroup may be adjusted independently of the others.

For example, with the 'Supernotch' (Q10 bands 3 & 4) in, its frequency may be altered by selecting the frequency

value windows of both bands 3 & 4, clicking the mouse within one of these, and 'dragging' the frequency value to either side. The width of the Supernotch can similarly be varied by selecting the Q value windows of Q10 bands 3 & 4, clicking the mouse within one of these, then 'dragging' the Q value to either side.
The frequency of the high-pass filter (Q10 bands 5, 6 & 7) may similarly be altered by selecting the frequency value windows of these three bands, and then click-and-dragging the frequencies to either side from within one of the selected value windows.

In all select, click and drag operations, the operation is applied to both stereo channels provided that Strap mode is selected. If Strap mode is not selected, then the operation is only applied to the channel (left or right) selected at that time.

In the **Q10 EQtools** setup, the frequency of the low-pass filter (Q10 bands 8 to 10) may be altered by selecting the 3 frequency value windows *and the three Q value windows* of these 3 bands 8 to 10, and then click-and-dragging (to either side) the frequencies and Q's simultaneously from within one of the selected value windows.

Note that for the low-pass filter, dragging the frequency values alone will not give a flat response before cut-off as the frequency is varied.

Editing of setups by selection and pasting

Any selection of settings of buttons and value windows can be copied over to any other by selecting, pressing the 'c' key on the keyboard, select the destination and pressing the 'v' key on the keyboard.

For example, in the **Q10 EQtools**, you may wish to use two Supernotches at different frequencies, but to include no Baxandall controls. Select both bands 3 & 4, press 'c' on the keyboard, select both bands 1 and 2, press 'v' on the keyboard, and you've pasted a second Supernotch to bandgroup 1 & 2.

Select and paste lets you create your own setups for fine-tuning. For example, you may paste six bands of pseudostereo from six adjacent bands of one of the 10-band pseudostereo setups, such as **Q10 PS10band**, and combine it with two Baxandall EQ bands copied from **Q10 EQtools**. Then you might add a Superparametric presence peak copied from two Q10 bands of **Q10 PresenceA**. The resulting setup can then be saved for future reuse and modification.

Chapter 1 - Pseudostereo

About Pseudostereo

Pseudostereo effects are intended to simulate a spacious or spread stereo sound from mono source material, and can be used to enhance old mono material, as well as to spread and 'liven up' mono components such as vocals, mono ambiences or mono instrumental lines within a stereo multitrack mix.

All the pseudostereo effects work by panning different frequency components of the sound alternately to either side of the stereo stage. Pseudostereo effects differ in the precise way stereo position varies with frequency. Most of the setups are characterized by the number of pseudostereo frequency 'bands' employed - the number of alternations of stereo position between left and right as frequency is varied across the audio band.

All the pseudostereo setups are intended for use with mono source material and are designed to minimize tonal changes in the original mono sound by retaining a flat total energy response from the two stereo loudspeakers.

The sound of all pseudostereo effects may be auditioned in stereo in Preview or Playback modes (in host applications where these modes are relevant) either with mono files or with stereo files having identical left and right channels ('double mono' stereo files). However, be aware that because of the processing limitations of Sound Designer II, if you wish to process a mono file to a pseudostereo stereo file, you must start off with a double-mono stereo file, otherwise, the processed file will still be in mono. You may convert a mono file into a double-mono stereo file rapidly in Sound Designer 2 by selecting 'Save a Copy...' in the file menu, then selecting Sound Designer II, 16-bit, stereo mode, entering a new file name, and then clicking the 'Save' button.

If used on stereo mixes, then the pseudostereo effects will work properly only for sounds in the center of the mix, with a risk of unpleasant tonal alteration to sounds panned to either side of the stereo stage. However, the effect may sometimes work well applied to stereo mixes where all the important sound sources are near the center, and where the stereo components are mainly reverb ambiences or spread sounds such as crowd noise, stereo rain sounds, or ADT'd sounds where the direct and delayed sound is applied at similar levels to opposite channels.

How well this will work for a specific piece of program material can only be determined by listening. When doing so, check for possible losses of frequency components in sounds at the two sides. The chances of success when applied to stereo mixes is increased if the 'narrow' versions of pseudostereo setups are used. However, these will provide a less dramatic spread of central mono sounds in the mix.

The principles of pseudo stereo are simple. We've already stated that different audio frequencies are panned alternately to the left and the right, but while the principle is simple enough, adjustments need to be very precise; even a very small inappropriate adjustment can ruin the general tonal balance. Unless you have plenty of time, and the necessary experience, to do fine tuning, it is not recommended that the tonal controls of the pseudostereo effects be adjusted individually, but that adjustments be confined to the select-and-drag options recommended in the following. These provide a wide range of adjustment in the pseudostereo effect while retaining the best available tonal neutrality.

It is important that all adjustments of pseudostereo setups should be made with the stereo channels strapped. If the Strap button is not lit up, click on it.

There are many basic varieties of pseudostereo effect available, each of which is also provided in a stereo reversed version.

Subjective aspects of pseudostereo

Pseudostereo effects have historically gained a very bad reputation, due to the very poor sound of most early attempts at this effect. Though the pseudostereo effects provided as setups, and their modifications as described below, are designed for maximum tonal neutrality, some subjective tonal effects are inevitable whenever the stereo image of a signal is manipulated. However, Q10 is able to handle pseudostereo processing with far fewer side effects than traditional systems.

Phasiness. Because the Q10 is a parametric equalizer with conventional so-called 'minimum phase' behavior, any non-flat setting causes some phase shifts as well as the frequency dependent amplitude gain variations seen on the display graph. When creating pseudostereo effects, this inevitably means that there will be phase differences between the stereo channels, which can sometimes be audible as an unpleasant 'phasiness'. The narrow setups will generally be found to be less phasey, and phasiness effects may also be found to be less noticeable in the 2-band setups.

The wide range of setups enables the user to select whichever version works best with the source material being used.

Mono Compatibility. Most of the pseudostereo setups provided are broadly mono-compatible, without suckouts or wide excursions in frequency response, although some setups will not give a completely flat frequency response when heard later in mono. Generally speaking, the historic, 10-, 14- and 18-band setups are fully mono compatible, and the tunable, 7-band and 10-bandX setups are generally the least mono compatible, but still usually acceptable. If mono compatibility is important, the 'narrow' versions of the setups will generally be more mono compatible than the normal. In all cases, a very wide pseudostereo sound can be obtained by clicking the polarity reverse button on the right channel only, but be aware that this increases phasiness and may lead to poor mono compatibility and low mono level. Again, this may be acceptable for background 'ambiences' accompanying a main dialog or vocal.

Headphone Results. It is very important to be aware that the subjective effect of stereo is very different over headphones and over loudspeakers, and it should not be assumed that a desired pseudostereo result over loudspeakers will automatically give a similar desired effect over headphones. If headphone listening is an important market for a mix, (e.g. on Walkmans), the results should be checked on both loudspeakers and headphones. This being said, the headphone results are usually found to be very acceptable, adding greatly to a sense of spaciousness.

Stereo Monitoring. While on the subject of monitoring, it is wise to check that the monitoring loudspeakers used are set up so as to give a good stereo imaging effect if subjective judgements on the pseudostereo are to be reliable. It is particularly important to listen in a precisely central symmetrical listening position to assess pseudostereo effects properly.

Gain settings and Overload. As provided, the input and output gain settings of the pseudostereo setups are such that the overall energy gain from input to output is unity, which permits a direct comparison between the bypass mode and the pseudostereo mode without gain changes. However, because the pseudostereo effects pan some frequency components mainly to one channel rather than splitting equally between two, this can sometimes result in channel overloads on the output in that channel. Usually, overloads can be avoided by pulling down both output gain faders - a further gain reduction of 3 dB will almost always avoid any overload problems, although the amount needed must be determined on a case-by-case basis.

Narrow and Reverse Setups. The reason for providing narrow, reverse and reverse narrow versions of setups is that these variations are often useful. For example, if a dominant frequency component is panned to the wrong side of the stereo stage, using the reverse version of the setup will pull it over to the other side. Also, if two mono sounds are both pseudostereo processed with the same setup, the result will simply be the same as the pseudostereo effect applied to a mono mix of the two, whereas if one is processed with a reversed pseudostereo setup to the other, the 'diversity' of a true stereo effect will be largely retained.

It is a good idea when processing several sounds in a mix with pseudostereo effects not to use identical pseudostereo processing on all, but to use reverse setups and different numbers of bands to create more variety of effect. Narrow setups are particularly recommended for material where an exaggerated spread should be avoided, notably speech. Narrow setups also tend to give less 'phasiness' and better mono compatibility.

The basic varieties are:

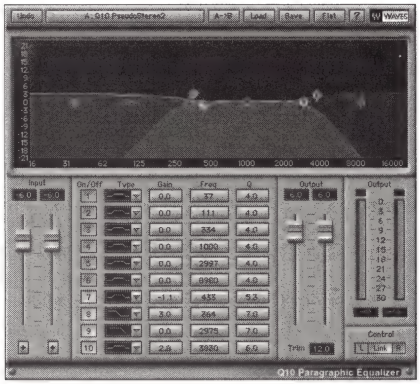

Q10 PseudoStereo2
Q10 PS2reverse (not pictured)
Q10 PS2band2 (not pictured)
Q10 PS2reverse2 (not pictured)

2-band pseudostereo effects, panning bass and treble to opposite sides of the stereo stage. A 6-band conventional graphic equalizer is also provided.

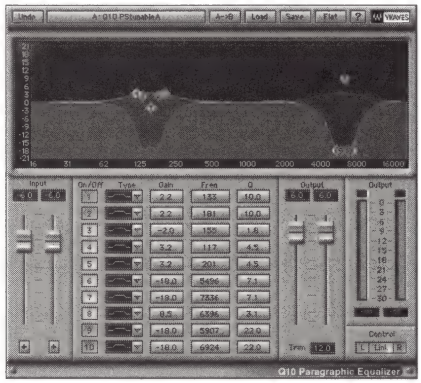

Q10 PStuneableA
Q10 PStuneableAreverse (not pictured)
Q10 PStuneableB (not pictured)
Q10 PStuneableBreverse (not pictured)
Q10 PStuneableC (not pictured)
Q10 PStuneableCreverse (not pictured)

These six are very highly tunable pseudostereo setups where individual bands can be panned to the left or the right to customize the stereo panning of different sound components.

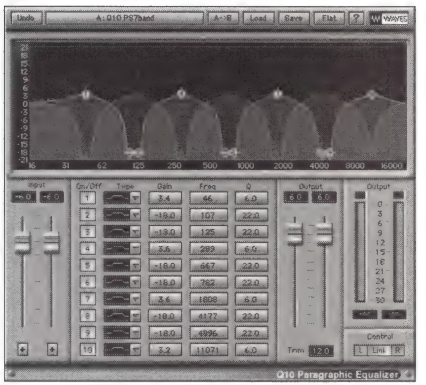

Q10 PS7band
Q10 PS7reverse (not pictured)
Q10 PS7band2 (not pictured)
Q10 PS7reverse2 (not pictured)

A 7-band pseudostereo effect, panning each of seven bands alternately to the left and to the right of the stereo stage. This effect is very wide, but not entirely mono-compatible. A highly adjustable effect.

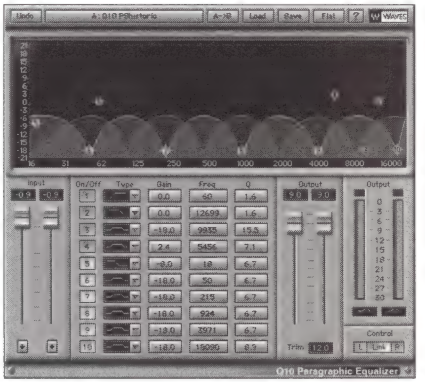

Q10 PShistoric
Q10 PShistoricreverse (not pictured)

A 10-band pseudostereo effect, panning each of 10 bands alternately to the left and to the right of the stereo stage. This setup is not adjustable in frequency and has good mono compatibility. The result is a diffused sound with 'normal' and 'narrow' spread versions. Four bands are reserved for normal EQ effects, preset as hiss and rumble filters, making this an ideal setup for reprocessing historic recordings.

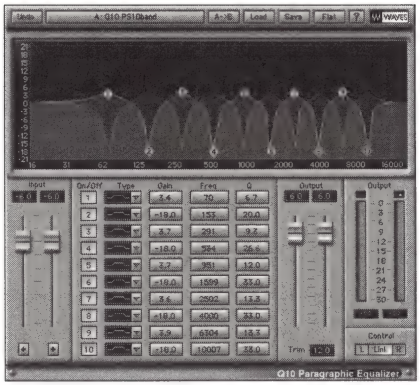

Q10 PS10band
Q10 PS10reverse (not pictured)
Q10 PS10narrow (not pictured)
Q10 PS10revnarrow (not pictured)

A 10-band pseudostereo effect, panning each of 10 bands alternately to the left and right of the stereo stage. Adjustable in frequency, giving a diffused sound, with 'normal' and 'narrow' spread versions. This setup has good mono compatibility.

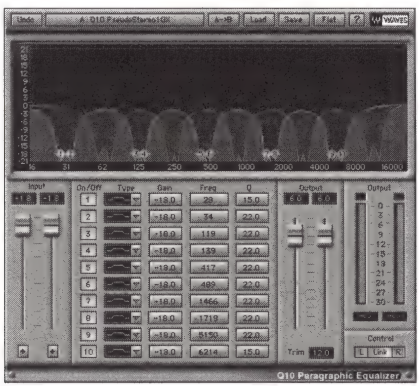

Q10 PseudoStereo10X
Q10 PS10Xreverse (not pictured)

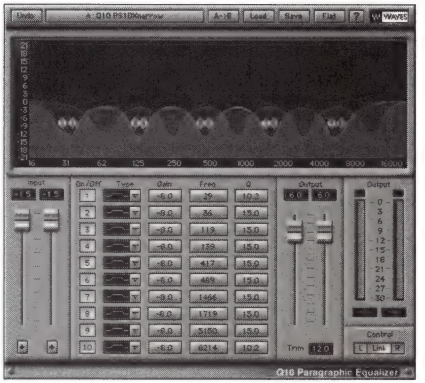

Q10 PS10Xnarrow
Q10 PS10Xrevnarrow (not pictured)

A 10-band pseudostereo effect, panning each of 10 bands alternately to the left and to the right of the stereo stage. This produces a very wide effect at the cost of poor mono-compatibility. This is possibly the most spectacular pseudostereo setup in the library. Not adjustable in frequency.

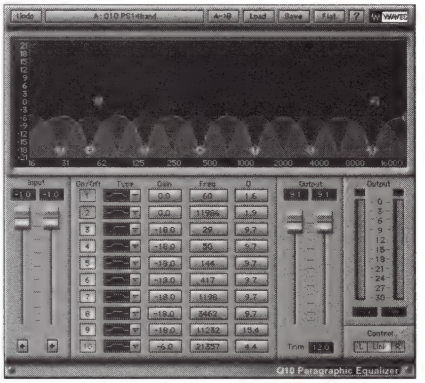

Q10 PS14band
Q10 PS14reverse (not pictured)

A 14-band pseudostereo effect, panning each of 14 bands alternately to the left and to the right of the stereo stage. This gives a diffuse pseudostereo effect, and reasonable mono compatibility, but is not spectacularly wide and tends to sound rather more colored than other effects. Not adjustable in frequency. A Baxandall equalizer is also provided.

Q10 PseudoStereo18
Q10 PS18reverse (not pictured)

An 18-band pseudostereo effect, panning each of 18 bands alternately to the left and to the right of the stereo stage. This gives the most diffuse pseudostereo effect, and reasonable mono compatibility, but is not spectacularly wide and again tends to sound rather more colored than other effects.

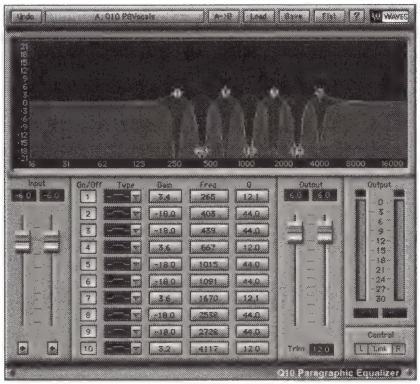

Q10 PSVocals
An adjustable pseudostereo effect where the spread is concentrated in the vocal frequency band.

Variations on these effects

All these pseudostereo effects are provided in a 'normal' and a reversed version, which simply has the equalization on the opposite channel to the normal version, so that the left/right panning effect of different frequencies can be reversed in the stereo stage.

Some of the pseudostereo effects also are also provided with 'narrow' and 'reversed narrow' versions which are essentially the same as the above versions except that the degree of stereo spread effect is reduced, for applications where a more subtle effect is preferred. Narrow setups are particularly recommended for speech where the normal spread generally sounds unnaturally wide.
All ten Q10 bands are used as a 'bandgroup' in providing the pseudo stereo effect for the 7, 10 and 18 band varieties. The 2-band varieties uses bands 7 to 10 as a bandgroup, the 10X-band variety uses bands 5 to 10 as a bandgroup, and the 14 band variety uses band 3 to 10 as a pseudostereo bandgroup.

Varying all frequencies together

Be sure to ensure that the two stereo channels are strapped when the any adjustments are made to pseudostereo effects. If you make a mistake and make an irretrievable change to the parameters, simply load the setup again and start over.

All pseudostereo effects, with the exception of the historic, 10X, 14 and 18 band cases, may have all bands moved up and down together. The procedure is to select and drag all the 'frequency' value windows only of the bandgroup (i.e. all bands in the tunable, 7-, and 10-band designs, and bands 7-10 in the 2-band designs). In this way, the operating frequencies of the pseudostereo effect may be moved up or down, This gives a change in the pseudostereo effect without altering its overall tonal balance. This often allows fine tuning to match the effect to particular program material, so that the stereo stage and the positions of frequency components within it are well balanced.

Separate Control of Pseudostereo Bands

Additionally, the tunable, 2-, 7- and 10-band pseudostereo setups allow separate control of individual pseudostereo frequency bands. The details differ in these three cases, and so are described individually.

2-Band Cases. (Q10 PseudoStereo2, Q10 PS2reverse, Q10 PS2band2and Q10 PS2reverse2).

Ensure that the two stereo channels are strapped when the following adjustments are made.

Bands 1 to 6 are not used for the pseudostereo effect so may be used for additional conventional equalization. As configured in the setup file, this provides a 6-band graphic EQ with band centers from 40 Hz to 12 kHz, and band gains may be adjusted up or down from the provided 0 dB flat setting.

The pseudostereo in this case is provided by Q10 bands 7 to 10. To switch the pseudostereo effect in or out, select and click the 4 In/Out buttons of Q10 bands 7 to 10.

This pseudostereo effect permits independent adjustment of the frequency at which bass signals are panned to one side and at which treble signals are panned to the other. To adjust the bass frequency, select and drag the *frequency* value windows of Q10 bands 7 and 8 as desired. To adjust the treble frequency, select and drag the *frequency* value windows of Q10 bands 9 and 10.

Be aware, however, that the pseudostereo effect will remain tonally neutral only so long as the bass and the treble bands do not overlap too much. The graphic display will show you this. Correct pseudostereo operation will occur if and only if the two curves cross over at about a 0 dB level. If the cross-over point level is significantly below 0 dB, the two bands overlap too much and the tonal quality will be compromised.

The two types of 2-band effect differ in the steepness of crossover between left, center and right positions, with the 'band2' effect giving a less rapid crossover.

7-band case. (Q10 PS7band, Q10 PS7reverse, Q10 PS7band2, Q10 PS7reverse2, Q10 PSVocals).

This setup is recommended where the maximum width of stereo effect is desired, and where mono compatibility is not of major importance. Like the 10-bandX effect, this pseudostereo effect can sometimes show more undesirable side effects on some material than many of the others, so the subjective result should be monitored carefully.

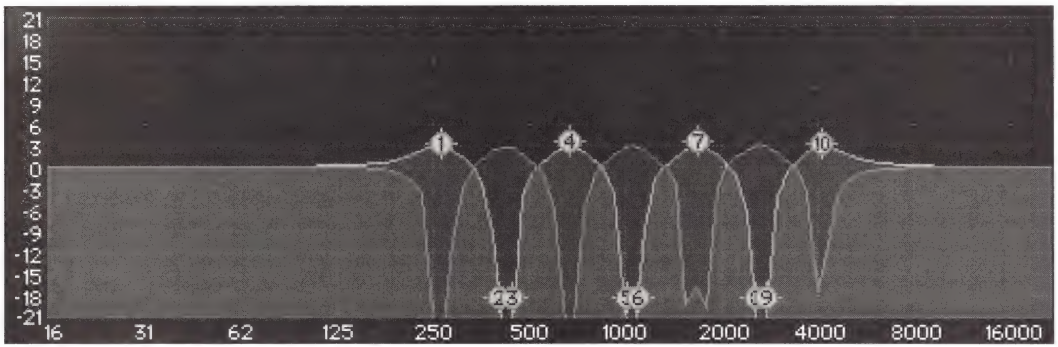

Frequency response of Q10 PSVocals

The '7band2' effect differs from the '7band' effect only in that the frequencies chosen for left and right panning are rather more widely spaced apart. By contrast, the PSVocals case has the bands packed together over a narrow vocal frequency range.

Once again, ensure that the two stereo channels are strapped when the following adjustments are made.

In this case, all ten Q10 bands are used for the pseudostereo effect. To switch the pseudostereo effect in or out, select the In/Out buttons of all ten Q10 bands, and then click anywhere within any of these buttons.

The frequencies of the pseudostereo effect may be adjusted up or down by select-and-dragging all ten frequency value windows in Q10 bands 1 to 10.

The 7-band versions may also be used as the basis for a more flexible adjustment of pseudostereo effect. The effect is divided into 4 bandgroups: bands 1-2, band 4-6, bands 7 -9 and band 10. Any number of these 4 bandgroups may be switched in or out separately, and their frequencies individually adjusted by selecting and dragging all frequency

value windows within a bandgroup, provided only that the bandgroups do not overlap too much (i.e. provided that the crossover of left and right graph curves does not fall significantly below 0 dB).

10-Band Cases. (Q10 PS10band, Q10 PS10narrow, Q10 PS10reverse, Q10 PS10revnarrow).

This 10-band pseudostereo effect is probably the most subtle and diffuse of the available setup types. It features a wide flexibility of adjustment and is recommended for when the frequency components simply require a broad spread and do not need to have pronounced stereo positioning.

Ensure that the two stereo channels are strapped when the following adjustments are made.
The pseudostereo in this case is provided by all ten Q10 bands. To switch the pseudostereo effect in or out, select and click all ten In/Out buttons of Q10 bands 1 to 10,.

The frequencies of the pseudostereo effect may be adjusted up or down by select-and-dragging all ten frequency value windows in Q10 bands 1 to 10.

In these setups, the spreading effect may be confined to limited parts of the frequency range only by switching out one or more bands. Any number of adjacent bands may be switched out at either end of the frequency range. In the middle of the frequency range, an even number of adjacent bands may be switched out. Remaining adjacent groups of switched in bands may be varied in frequency by selecting and dragging their frequency value windows. Tonal neutrality will be retained provided that the bands do not overlap too much (i.e. provided that the crossover of left and right graph curves does not fall significantly below 0 dB).

Tunable Pseudostereo Setups

Q10 PStuneableA
Q10 PStuneableAreverse
Q10 PStuneableB
Q10 PStuneableBreverse
Q10 PStuneableC
Q10 PStuneableCreverse

These pseudostereo setups allow tuning of particular frequency bands to the left side or the right side of the image, and are particularly effective for placing elements of the original mix at different positions within the stereo spread. These setups provide the means of doing this providing the sounds being processed are separated enough in frequency.

The four bandgroups for these presents are:

Q10 bands 1 & 2 : narrow frequency band panned to left (right for 'reverse' setups)
Q10 bands 3 to 5 : wide frequency band panned to left (right for 'reverse' setups)
Q10 bands 6 to 8 : wide frequency band panned to right (left for 'reverse' setups)
Q10 bands 9 & 10 : narrow frequency band panned to right (left for 'reverse' setups)

Any one or more bandgroups may be adjusted in frequency by selecting the frequency value windows of that

bandgroup(s), and dragging to either side.

The bandgroups may be switched in or out individually, provided that no two 'in' bandgroups overlap too much in frequency range. Excessive overlap can be seen from the graphical display by the crossover of left and right curves for adjacent bands differing (by more than say a dB) from 0 dB. It is normally a good idea for two adjacent panned bands to be to opposite sides of the stereo stage - since this allows much more overlap of bands than the case where they are panned the same way.

The A, B and C versions of the setups differ only in the chosen initial frequencies and in which bandgroups are initially switched in. Since, apart from the initial center frequencies, their bandgroups are identically set up, they can all be used with frequency tuning to achieve identical effects, but it is easiest to choose the version closest to what is required as a starting point.

The A versions of the tunable pseudostereo setups are ideal for achieving a 2-band pseudostereo effect, where one wishes to be able to select using a narrow or wide frequency band after tuning the pseudostereo panning in frequency. In the A version, only two bandgroups are switched in when loaded - the wide frequency bands. By selecting all 10 In/Out buttons and switching, the wide bands are switched out and narrow bands tuned to the same frequencies are switched in. The idea is to tune the frequencies of bands 1 to 5 together by selecting all their frequency value windows and dragging. Similarly, tune the frequencies of bands 6 to 10 together by selecting all their frequency value windows and dragging.

Having chosen the frequencies at which the desired sounds are panned to left and right, you can then choose whether to have narrow or wide frequency range for the panning by selecting the In/Out buttons of bands 1 to 5, and clicking on one of them. This will alternate between a narrow and a wide band at the same frequency. The same procedure can be used independently to select a narrow or wide frequency range for the panning in bands 6 to 10.

The B and C versions of the tunable pseudostereo setups are intended as a starting point for a 4-band tunable pseudostereo effect. They all have the same frequency bands panned to the same positions when loaded. The only difference is the order of narrow and wide bands.

In the B versions, the lowest and highest bands have narrow frequency range, and the mid-low and mid-high bands have wide frequency range. In the C versions, the lowest and highest bands have wide frequency range, and the mid-low and mid-high bands have narrow frequency range.

Inflexible Pseudostereo Setups

Q10 PseudoStereo10X
Q10 PS10Xreverse
Q10 PS10Xnarrow
Q10 PS10Xrevnarro
Q10 PShistoric
Q10 PShistoricreverse
Q10 PS14band
Q10 PS14reverse
Q10 PseudoStereo18
Q10 PS18reverse

These pseudostereo setups are inflexible in offering neither frequency adjustment nor spread adjustment, but in compensation offer a remarkably large number of pseudostereo bands, respectively, 10, 10, 14 and 18, for a much more diffuse effect than setups with a small number of bands. Because there are more bands, there's less likelihood of individual sounds having different frequency components 'split' between different stereo positions.

The '10X' pseudostereo setups provide a 10 band pseudostereo with a very wide stereo stage for when spectacular spread is desired without excessive phasiness. However, mono compatibility is not particularly good.

The 18 band setups provide a very diffuse sound, but not the widest stage, but have fairly good mono compatibility making them suitable for use in areas such as broadcast when a general-purpose pseudostereo effect is required.

The 14-band setups are similar to the 18-band case, except that they only have 14 pseudostereo bands. However, the pseudostereo algorithm now uses as bandgroup Q10 bands 3 to 10, which leaves Q10 bands 1 and 2 free for conventional equalization. As loaded, bands 1 and 2 are configured as Baxandall bass and treble controls, and adjustment of their gains provides respective adjustment of bass and treble. This setup is useful in applications where it is desirable to accomplish EQ and pseudostereo processing in one pass.

The 'historic' setups are similar to the 14- and 18-band pseudostereo, except they have only 10 pseudostereo bands, implemented by Q10 bands 5 to 10. This leaves Q10 bands 1 to 4 free for conventional equalization, when the power of a 4-band parametric EQ is needed. As loaded, band 1 is configured as a high-pass filter for filtering out bass rumble, and Q10 bands 2 to 4 form a bandgroup for a tunable psychoacoustic hiss filter. By selecting the frequency value windows of bands 2 to 4 and dragging, the frequency of action of the hiss filter may be varied.

As loaded, this setup is suitable for the treatment of historic mono material, both to achieve a pseudostereo effect, and to filter hiss (or 78 rpm record scratch noise) and rumble in one EQ pass. The hiss filter has been designed for minimum subjective tonal degradation for a given degree of hiss reduction.

HINT. In the historic and 14-band setups, you can see the EQ effect on the graphic display by selecting the bandgroup In/Out buttons of the pseudostereo bandgroup and clicking on one of them. This will remove the pseudostereo and allow you to see the EQ effect. Clicking again will bring in the pseudostereo effect again.

Chapter 2 - Crossover

About crossover

The 'crossover' setups are all designed to split incoming mono signals into two sharply separated bands. Typically, one is in the left output channel, the other in the right where the crossover slope between the bands is of the order of 50 dB per octave. The crossover points are at -3dB so that the total energy in the left and the right channel is roughly flat.

The crossovers are designed so that simply selecting all ten frequency value windows & dragging moves the crossover frequency or frequencies up and down.

These crossover setups can be used to create extreme form of 'pseudostereo' where everything in one frequency band is on the left, everything in another frequency band is on the right, and everything in yet another frequency band may be in the middle. With crossover frequencies matched to the musical content, this can place instruments occupying different frequency bands in any of three positions.

The mono compatibility of the crossover effect cannot be guaranteed because of response irregularities in the crossover region due to phase addition and subtraction. However, in cases where the crossover is very sharp, this will only affect a very narrow frequency range, which in practice may be acceptable.

Note that all adjustments of 'crossover' effects must use strapped channels. If a file is to be processed, the file should be a mono signal copied as a stereo file, since stereo output files can only be generated from stereo input files.

Q10 crossover1

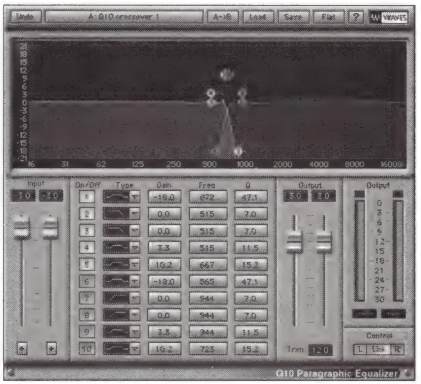

Q10 crossover1

This setup simply splits the incoming signal placing a brickwall low-pass filter in the left channel and a brickwall high-pass filter in the right channel. This setup uses as bandgroup Q10 bands 1-5, placing bass on the left and treble on the right. As loaded, the crossover frequency is 700 Hz, but it may be select and dragged to any other frequency up to several kHz.

Bandgroup Q10 channels 6 to 10 provides an identical crossover except that treble is sent to the right and bass to the left. One may switch between the two crossovers by selecting all ten In/Out buttons and clicking on one of them.

Q10 Crossover 2, 3, and 4

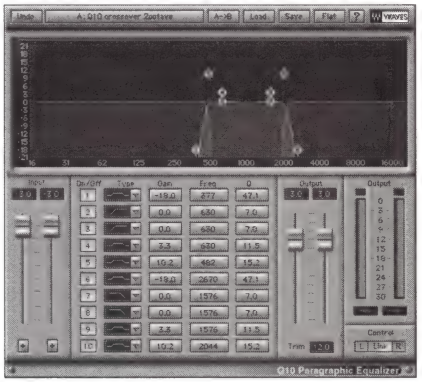

Q10 crossover2octave

Q10 crossover3octave

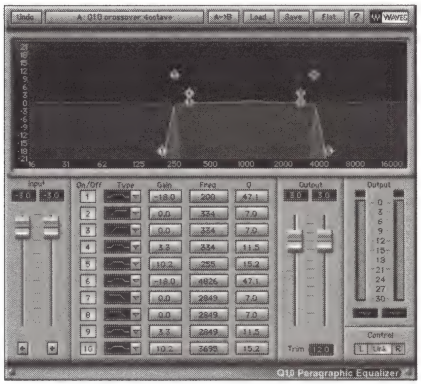

Q10 crossover4octave

These use all 10 Q10 bands as a bandgroup, and put low and high frequencies on the right channel and middle frequencies on the left. The two crossover frequencies are about 2, 3 and 4 octaves apart in the respective setups, with the middle band centered at 1 kHz. The overall frequency may be moved up or down by selecting all ten frequencies and dragging.

A practical example is to use, say, the crossover4octave case to put the vocals of a mono recording on the left, and the bass line and cymbals on the right.

Q10 crossoverLCR

Q10 crossoverLCR2

These setups put low frequencies below a first low crossover frequency on the left, high frequencies above a second high crossover frequency on the right, and everything between the two crossover frequencies in the center at equal levels in both channels. The LCR setup has the spacing between crossover frequencies of 4 octaves, and the LCR2 setup uses 3 octave spacing.

The response of each channel steps down 3 dB over the middle band to ensure that the total energy of the 2 channels remains constant.

The first low crossover frequency is controlled by bandgroup Q10 bands 1 to 5, and may be adjusted up or down by selecting all five frequency value windows of this bandgroup and dragging them. The second high crossover frequency is controlled by bandgroup Q10 bands 6 to 10, and may be adjusted up or down by selecting all 5 frequency value windows of this bandgroup and dragging them.
Please note that at middle frequencies, the phase response on the two channels is not identical, so that there will be some 'phasiness' and imperfect mono compatibility.

CR/LC-types

Q10 crossoverCR
Q10 crossoverLC (not pictured)

These two setups are identical except that the channels are the opposite way round.

The LC setup when opened has bandgroup Q10 bands 1 to 5 in, and bandgroup Q10 bands 6 to 10 out. It has the effect of putting frequencies below a first crossover frequency on the left, and frequencies above that crossover frequency in the center. The crossover frequency may be moved up or down by selecting the 5 frequency value windows of Q10 bands 1 to 5 and dragging.

If all 10 In/Out buttons are selected and clicked, so that Q10 bands 6 to 10 are in and bands 1 to 5 are out, then the LC setup has the effect of putting frequencies above a second crossover frequency on the left, and frequencies below that crossover frequency in the center. The crossover frequency may be moved up or down by selecting the 5 frequency value windows of Q10 bands 6 to 10 and dragging.

The CR setup is identical except that panning to the left is replaced by panning to the right throughout.

Other uses of crossover setups

The Crossover setups may also be used as conventional active crossovers for use in the design of loudspeaker systems or in any other test situation where a versatile, precision crossover is required. When used in TDM systems or within WaveShell, signals may be processed in real time without first having to be stored as soundfiles.

Rapid crossover types

The following "crossover" setups for the Q10 split the audio band into several bands and feed them to left and right channels. They are characterized by flat responses within each frequency range and very rapid crossovers between frequency bands.

Q10 Xover2L
Q10 Xover2R
Q10 Xover2Low
Q10 Xover2Hi

An ultra-rapid crossover bandsplit into 2 bands where the crossover rate is around 150 dB per octave down to about -40 dB. These setups use bandgroup 1-8, and the cross-over frequency may be adjusted (up to a few kHz) by selecting all 8 frequencies in both channels, and dragging up or down.

Q10 Xover2L puts the low-pass filter in the right channel and the high-pass in the left.

Q10 Xover2R puts the low-pass filter in the left channel and the high-pass in the right.

Q10 Xover2Lowputs the low-pass filter in both channels.

Q10 Xover2Hiputs the high-pass filter in both channels.

The use of these setups allows an effectively clean split of any audio signal into two separate frequency bands, which can be separately processed before being reassembled. Only in the narrow crossover frequency region will there be any notable side effects and interaction between bands.

Q10 Xover3midlow
Q10 Xover3midhi
Q10 Xover3low
Q10 Xover3mid
Q10 Xover3hi

A rapid crossover bandsplit into 3 bands where the crossover rate for the low and high pass is around 150 dB per octave down to about -40 dB, and where the mid-passs has rather gentler cross-over rate. The low and high-pass filters use bandgroup 1-8, and the mid-pass filter uses bandgroup 1-5 for its high-pass action and bandgroup 6-10 for its low-pass action.

As provided, the mid-pass band covers about 4 octaves, and the cross-over frequencies may be adjusted (up to a few kHz) by selecting all 10 frequencies on both channels and dragging up or down.

Q10 Xover3midlow puts the low-pass filter in the right channel and the mid-pass in the left.

Q10 Xover3midhi puts the mid-pass filter in the left channel and the high-pass in the right.

Q10 Xover3octlowputs the low-pass filter in both channels.

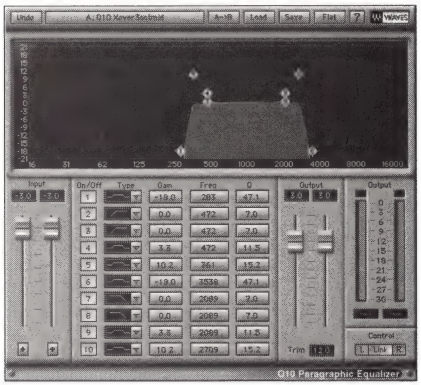

Q10 Xover3octmidputs the mid-pass filter in both channels.

Q10 Xover3octhputs the high-pass filter in both channels.

The use of these setups allows a reasonably clean split of any audio signal into three separate frequency bands, which can be separately processed before being reassembled. Only in the crossover frequency regions will there be any notable side effects and interaction between bands.

LRLR types

Q10 XoverLRLR1.4oct
Q10 XoverRLRL1.4oct
Q10 XoverLRLR1.9oct
Q10 XoverRLRL1.9oct
Q10 XoverLRLR2.4oct
Q10 XoverRLRL2.4oct

A rapid crossover bandsplit into 4 bands where the rejection out of band exceeds -35 dB in the low-mid and high-mid bands. One channel passes the low and high-mid bands and the other the low-mid and high bands. The width of the low-mid and high-mid bands are equal, and are respectively 1.4, 1.9 and 2.4 octaves depending on the setup chosen.

All 10 bands are used, and the cross-over frequencies may be moved up or down together, up to a few kHz, by selecting and dragging all ten frequencies in both channels.

Q10 XoverLRLR1.4oct

Q10 XoverRLRL1.4oct

The LRLR setups put the low and high-mid filtering in the left channel and the low-mid and high filtering in the right channel. The RLRL setups put the low and high-mid filtering in the right channel and the low-mid and high filtering in the left channel.

With a dual-mono input, these setups may be used to obtain a fairly extreme pseudostereo effect with a pronounced left-right split.

3-way types

Q10 Xover30ctCLR
Q10 Xover30ctCRL
Q10 Xover3.50ctCLR
Q10 Xover3.50ctCRL
Q10 Xover40ctCLR
Q10 Xover40ctCRL

A rapid crossover bandsplit into 3 bands designed primarily for pseudostereo processing of dual-mono inputs, where the low frequency band is assigned to stereo center, the mid band to one channel and the high band to the other. The separation in the mid band region exceeds 35 dB.

The setups use all 10 bands, and all crossover frequencies may be moved up or down together, up to a few kHz, by selecting and dragging all ten frequencies in both channels. The width of the mid band is 3, 3.5 or 4 octaves depending on the setup chosen.

Q10 Xover30ctCLR

Q10 Xover30ctCRL

The CLR setups put the mid filtering in the left channel and the high filtering in the right channel. The CRL setups put the mid filtering in the right channel and the high filtering in the left channel.

Q10 PlateauCLR-X
Q10 PlateauCRL-X

These two setups have a similar rapid 3-way pseudostereo bandsplit to the previous setups, but with different performance specifications and adjustability.

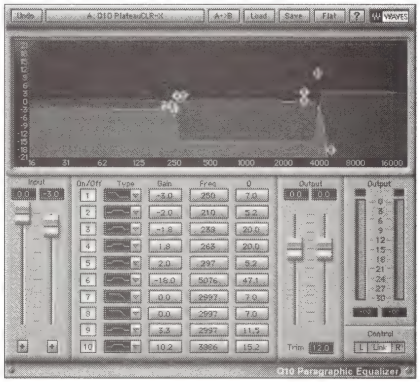

Q10 PlateauCLR-X

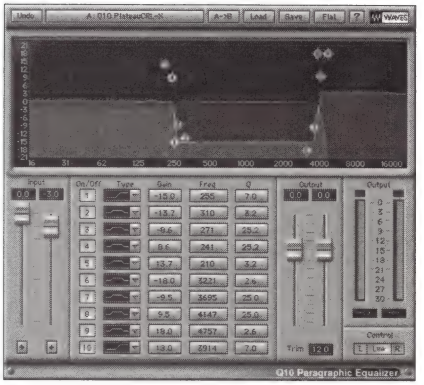

Q10 PlateauCRL-X

In these setups, the mid-band separation is only 18 dB, which gives a noticably narrowed pseudostereo effect. The high-band separation is still wide.

However, the crossover frequencies may be adjusted independently up to a few kHz. Bandgroup 1-5 controls the lower cross-over frequency and bandgroup 6-10 the higher crossover frequency. A crossover frequency may be moved up or down, up to a few kHz, by selecting and dragging all five frequencies in both channels in the selected bandgroup.

Chapter 3 - Bandlimiting

Many audio signals are destined to be heard via systems or transmission channels that impose a quality limitation on the material being transmitted, and it is often helpful to know what audible effects such a system will have on a given piece of program material. For example, if important parts of a mix become inaudible or out of balance after being passed through a Q10 emulation of the transmission system, corrective action (such as preparing a specially EQ'd or balanced mix) may be taken. Three of the setups provided are particularly useful in this type of application and are very similar except for the precise choices of cut-off frequencies in the bass and treble:

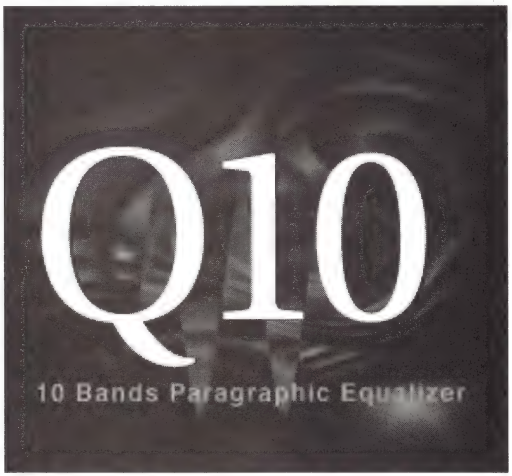

Q10 telephone: simulates a typical telephone bandwidth of 300 Hz to 3 kHz, using a very steep 'brickwall' cut-off outside this band. Unlike many typical telephones, however, it is flat within this band. This setup is useful both for determining the effect of program material likely to be heard down telephone lines and for simulating a 'telephone' effect when needed for drama, commercials or an effect in music.

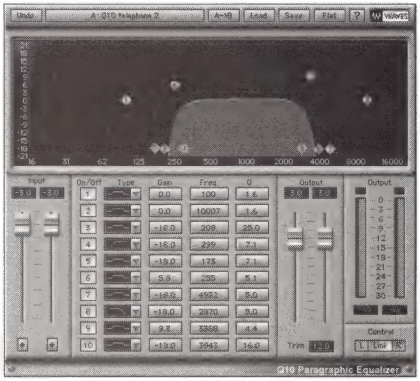

Q10 telephone2: also simulates a typical telephone bandwidth of 300 Hz to 3 kHz, using a rather less steep cut-off outside this band. Again the response is flat within this band. Q10 bands 1 and 2 allow extra equalization to be provided but are initially set to flat. If their gains are altered, they provide Baxandall bass and treble controls.

Q10 AMradio: simulates a typical received AM radio bandwidth of 100 Hz to 5 kHz, using a very steep cut-off outside this band. Unlike many typical AM radios, however, it is flat within this band. **Q10 AMradio** provides the steepest brickwall filtering for the best effect. This setup is useful both for determining the effect of program material likely to be heard on AM radios and for simulating a 'radio' effect when needed for drama, commercials or as an effect in music.

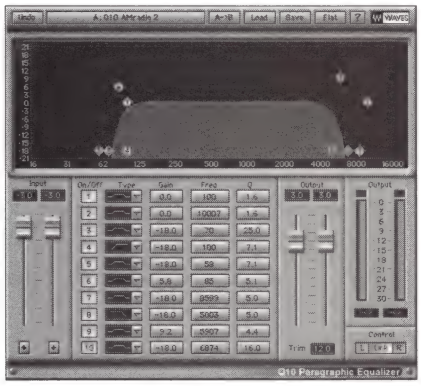

Q10 AMradio2: provides a less sharp cutoff simulation of the AM radio bandwidth 100 Hz to 5 kHz. Q10 bands 1 and 2 allow extra equalization to be provided but are set to flat as provided. If their gains are altered, they provide Baxandall bass and treble controls.

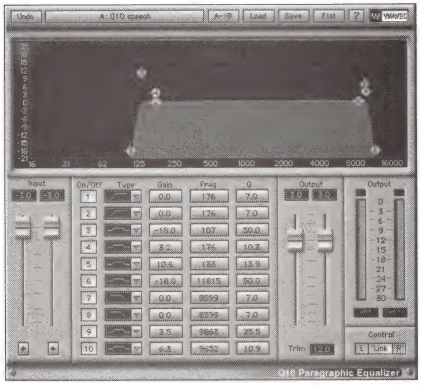

Q10 speech: provides very steep 'brickwall' cut-offs at 130 Hz and 10 kHz, and a flat response in between. It is designed for limiting the bandwidth of speech in applications such as multimedia or some broadcast applications where an extended speech bandwidth may be undesirable. Typically, this setup may be used when speech intelligibility and clarity is more important than true fidelity, or where the loudspeaker system on which the material will be reproduced has a limited bass response.

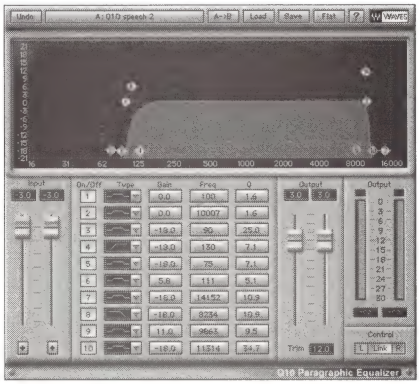

Q10 speech2: provides slightly less sharp cut-offs at 130 Hz and 10 kHz, and a flat response in between. Q10 bands 1 and 2 allow extra equalization to be provided and are set to flat as provided. If their gains are altered, they provide Baxandall bass and treble controls.

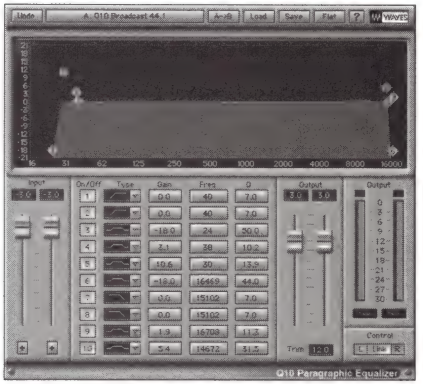

Q10 Broadcast44.1

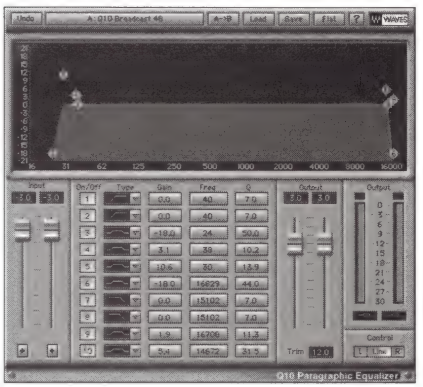

Q10 Broadcast48

Q10 Broadcast44.1, Q10 Broadcast48. permits files recorded at the respective 44.1 and 48 kHz sampling rates to be bandlimiting to from 30 Hz to 15 kHz. This is the frequency range used for FM broadcasting and TV work. This setup can be used as an antialiasing filter prior to sampling rate conversion to the widely used broadcast sampling rate of 32 kHz. Not adjustable.

Adjusting these setups

Q10 telephone, Q10 AMradio, Q10 speech

These three setups all allocate the Q10 bandgroup bands 1 to 5 for the high-pass (low cut) filter, and the Q10 bandgroup bands 6 to 10 for the low-pass (treble cut) filter.

The high-pass filter and low-pass filter may be used individually or together.

High-pass filter. This uses the bandgroup Q10 bands 1 to 5. To switch the high-pass in or out, select thpcì¨yn/Out buttons of Q10 bands 1 to 5, and click on any one of these buttons. To adjust the frequency of cutoff of the high-pass filter up or down, select the 5 frequency value windows of Q10 bands 1 to 5, and click and drag the frequency on any one of them.

Low-pass filter. This uses the bandgroup Q10 bands 6 to 10. To switch the low-pass in or out, select the 5 In/Out buttons of Q10 bands 6 to 10, and click on any one of these buttons. To adjust the frequency of cutoff of the high-pass filter up or down, select the 5 frequency value windows of Q10 bands 6 to 10, and click and drag the frequency on any one of them.

Note from the graphical display that the frequency range over which the brickwall low-pass remains reasonably flat varies with the cut-off frequency setting, and for best results it is advisable to start with the setup with the low-pass cut-off frequency closest to the one desired.

Q10 telephone2, Q10 AMradio2, Q10 speech2

These three setups all allocate Q10 bands 1 and 2 to the respective bass and treble of a Baxandall tone control, and allocate the Q10 bandgroup bands 3 to 6 for the high-pass (low cut) filter, and the Q10 bandgroup bands 7 to 10 for the low-pass (treble cut) filter.

The high-pass filter and low-pass filter may be used individually or together.

High-pass filter. This uses the bandgroup Q10 bands 3 to 6. To switch the high-pass in or out, select the four In/Out buttons of Q10 bands 3 to 6, and click on any one of these buttons. To adjust the frequency of cutoff of the high-pass filter up or down, select the four frequency value windows of Q10 bands 3 to 6, then click and drag the frequency on any one of them.

Low-pass filter. This uses the bandgroup Q10 bands 7 to 10. To switch the low-pass in or out, select the four In/Out buttons of Q10 bands 7 to 10, and click on any one of these buttons. To adjust the frequency of cutoff of the high-pass filter up or down, select the four frequency value windows *and the four Q value windows* of Q10 bands 7 to 10, then click and drag the frequency and Q together on any one of them. Dragging frequency alone will not give a flat response before cut-off as the frequency is varied. In practice, the low-pass filter will work well for '-3 dB' cut-off frequencies in the range 2 to 13 kHz approximately.

Alternative low- and high-pass filters.

Alternative low- and high-pass filters with a slightly gentler cut off rate are provided in **Q10 EQtools**. In this instance, the high pass filter uses the bandgroup of Q10 bands 5 to 7, and may be switched in and out by selecting and clicking the three buttons of this bandgroup, and the cut-off frequency may be varied by selecting and dragging the three frequency value windows of this bandgroup.

In this case, the low-pass filter uses the bandgroup of Q10 bands 8 to 10, and may be switched in and out by selecting and clicking the three buttons of this bandgroup, and the cut-off frequency may be varied by selecting and dragging the three frequency value windows *and the three Q value windows* of this bandgroup. For the low-pass filter, dragging frequency alone will not give a flat response before cut-off as the frequency is varied.

Hiss filtering. The above low-pass filters can be used to reduce hiss on hissy material, but the user is cautioned that this is generally not the most effective form of filtering for reducing audible hiss with minimum side effect. This is because the ears are most sensitive to the frequency region around 4 kHz, with a low sensitivity to very high frequencies. A low-pass filter removes all the very high frequencies, despite the fact that the ears are less sensitive to hiss at these highest frequencies. A better strategy is to use filters designed to reduce or remove only those frequencies at which hiss is most audible, and to leave lower *and higher* frequencies alone. Such filters are provided in other setups as described below.

Also, very steep-cut, low-pass filters are found to give the remaining hiss an unpleasant 'bright' and edgy quality that is sometimes more objectionable than leaving hiss in. The use of low-pass filters to remove hiss does not generally produce the best subjective results.

Superparametric EQ

Parametric equalizers, while extremely useful, rarely provide exactly the subjective effect required by audio production engineers. Their action might more appropriately be described as approximate, because the ears are surprisingly sensitive to small variations in frequency response shape, Even relatively small boosts or cuts away from the desired frequencies can be unpleasantly audible.

Conventional parametric 'bell' boosts and cuts not only boost and cut the desired frequency band, but they also affect frequencies well away from the target frequencies causing unpleasant coloration.

Several setups provided here provide 'Superparametric' boosts and cuts that are more precise in their action, affecting only the desired frequency range and having virtually no effect outside from that range. When used to cut a desired frequency band, the result is less tonal dulling of other frequencies, and when used to boost a desired frequency band, the result is less tonal alteration of the rest of the sound.

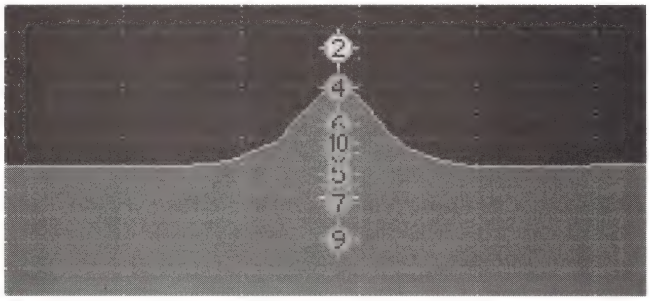

Frequency response of Superparametric boost.

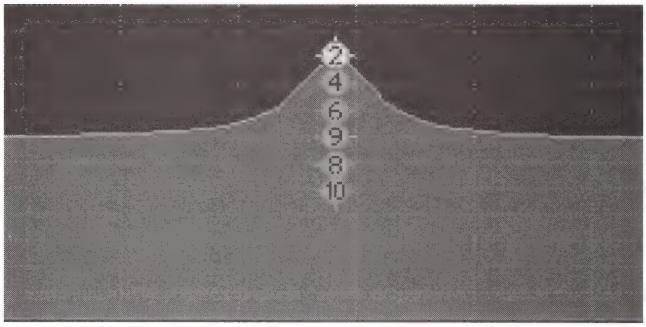

Frequency response of ordinary parametric boost.

To hear the difference between parametric and Superparametric boosts, load **Q10presenceA** in setup A, and load **Q10presenceB** in setup B. **Q10presenceA** is a 9.3 dB Superparametric boost at 3.5 kHz, whereas **Q10presenceB** provides an ordinary 9.3 dB parametric boost at 3.5 kHz, carefully matched in the boost band. Switching between setup A (Superparametric), setupB (parametric) and bypass on speech, it will be heard that the Superparametric EQ is tonally more natural than the parametric.

Adjusting Superparametric EQ

Q10 PresenceAprovides five alternative degrees of boost or cut at the presence frequency of 3.5 kHz.

The bandgroup Q10 bands 1&2 provide +6.2 dB boost at 3.5 kHz.
The bandgroup Q10 bands 3&4 provide +9.3 dB boost at 3.5 kHz.
The bandgroup Q10 bands 5&6 provide +3.1 dB boost at 3.5 kHz.
The bandgroup Q10 bands 7&8 provide 3.1 dB cut at 3.5 kHz.
The bandgroup Q10 bands 9&10 provide 6.2 dB cut at 3.5 kHz.

Select the In/Out button of the desired bandgroup to obtain the desired degree of boost or cut.

To vary the center frequency, select all frequency value windows and drag the frequency to either side as desired. The value windows show the selected frequency. To vary the width of the Superparametric EQ, select all Q value windows and drag the Q to either side as desired.

Q10 PresenceF: A shortcoming of the previous setup is the crude set of values of gain provided - only 6 values (including flat 0 dB) in 3.1 dB steps. If a finer and wider range of values is desired, Q10 PresenceF may be used instead, although selection of gain is much less straightforward. Here also, only a limited range of gains is provided, but other gains may be achieved by switching in two or more bandgroups. The gain may by altered in 2 dB steps between -12 dB and + 12 dB.

Five bandgroups provide five alternative degrees of boost or cut at the presence frequency of 3.5 kHz.

The bandgroup Q10 bands 1&2 provide +12 dB boost.
The bandgroup Q10 bands 3&4 provide +6 dB boost.
The bandgroup Q10 bands 5&6 provide -2 dB cut .
The bandgroup Q10 bands 7&8 provide -4 dB cut.
The bandgroup Q10 bands 9&10 provide -6 dB cut.

To obtain the desired degree of boost or cut, switch in the following bandgroups:

+12 dB, bandgroup 1&2 only.
+10 dB, bandgroups 1&2 and 5&6 only. (10 dB = 12 dB - 2 dB)
+8 dB, bandgroups 1&2 and 7&8 only. (8 dB = 12 dB - 4 dB)
+6 dB, bandgroup 3&4 only.
+4 dB, bandgroups 3&4 and 5&6 only. (4 dB = 6 dB - 2 dB)
+2 dB, bandgroups 3&4 and 7&8 only. (2 dB = 6 dB - 4 dB)
0 dB, all bandgroups out.
-2 dB, bandgroup 5&6 only.
-4 dB, bandgroup 7&8 only.
-6 dB, bandgroup 9&10 only.
-8 dB, bandgroups 5&6 and 9&10 only. (-8 dB = -6 dB - 2 dB)
-10 dB, bandgroups 7&8 and 9&10 only. (-10 dB = -6 dB - 4 dB)
-12 dB, bandgroups 5&6, 7&8 and 9&10 only. (-12 dB = -6 dB - 4dB - 2 dB).

To vary the center frequency, select all frequency value windows and drag the frequency to either side as desired - the value windows show the selected frequency. To vary the width of the Superparametric EQ, select all Q value windows and drag the Q to either side as desired.

Uses of Superparametric EQ

The Superparametric EQ's ability to operate on a selected band of frequencies without affecting neighboring frequencies has a variety of different uses.

One application is to make a sound 'buried' in a mix more audible, by selectively boosting its most prominent frequencies. For example, buried vocals in live recordings can often be 'pulled out' by boosting frequencies between 1.5 or 2 kHz. Using Superparametric EQ, this can be done with less effect on the rest of the mix.

A converse application is to pull down an over-prominent sound in a mix by cutting its most prominent frequencies using a Superparametric cut, again with less likelihood of adversely affecting the rest of the mix.

A more sophisticated use is to improve a mix where several sounds have similar frequency balances causing them to get in each other's way. By selectively boosting some sounds and cutting others in the same frequency band using a Superparametric EQ (ideally before mixdown), it is possible to alter the relative audibility of individual sounds in the mix while having very little subjective effect on the overall sound of the mix.

To a limited extent, even applying a Superparametric boost in one stereo channel and a complementary cut on the other can sometimes rescue a marginal mix, if the sounds affected are towards different sides of the stereo stage. However, attempting to overcompensate using this method may lead to a noticeable deterioration in sound quality.

A further important use of Superparametric EQ is for filtering out unwanted noises, including rumble noise, noises due to air conditioning, and hiss. This can be used in much the same manner as parametric EQ, except that the effect on other sounds is lessened. A general hint here is to locate the problem frequency by setting up a Superparametric boost, with fairly narrow width, and to tune it in frequency until the problematic noise is most audible. A Superparametric cut may then be applied at this frequency. It is also helpful to reduce the width of the Superparametric notch so that it is as narrow as possible consistent with the desired degree of reduction of the unwanted noise. The optimum setting is the once which achieves the best compromise between the removal of unwanted noise and audible degradation of the wanted signal.

When a deep notch is required as a noise filter, you can use the Supernotch (Bandgroup Q10 bands 3 & 4) of the **Q10 EQtools**. This may be adjusted in frequency by selecting the two frequency value widows and dragging to either side as required, and may be adjusted in width by selecting the two Q value windows and dragging to either side.

Superparametrics used in this way can make very effective hiss filters. They are generally far more effective than using low-pass filters, however, one warning is in order. Some less good monitoring speakers may well have prominent emphasis frequency bands, and there is a risk with Superparametric dips or notches that you may be tuning out a loudspeaker resonance rather than a true peak in the audible noise. Use of good neutral monitoring loudspeakers will minimize this risk, but it is a wise precaution where possible to check the results of mixes done with Superparametric filters on alternative high quality monitoring situations.

Deep Supernotch

Q10 Supernotch

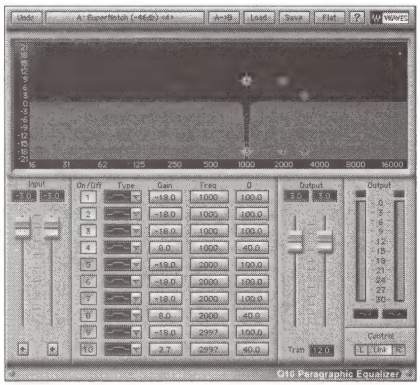

This type of filter is intended to remove whistles or highly pitched interference.

This setup provides a deep narrow tunable notch with about 46 dB attenuation, using as bandgroup Q10 bands 1 to 4. Its frequency may be tuned by selecting the four frequency value windows and dragging. The bandgroup Q10 bands 5 to 8 provides a second similar Supernotch tuned to twice the frequency, which may be switched in if there is a second harmonic present, or tuned to a second independent whistle frequency. The bandgroup Q10 bands 9 & 10 provides a 15 dB attenuation Supernotch set to three times the frequency for the case where a third harmonic component is also present, or where another frequency requiring less attenuation needs to be removed.

All frequencies may be dragged up and down together if all 10 frequency value windows are selected.

Psychoacoustic noise filters

While Superparametric notches are excellent for reducing noises that are most predominant around a discrete frequency, it is common to encounter material with broadband rumble or hiss noise where Superparametric notches are not adequately effective for reducing noise. In this situation, low-pass filters are needed to remove hiss and high-pass filters to reduce rumble. But psychoacoustically, steep cut filters are far from optimum in this application because they apply an unnecessary degree of filtering to less audible noise components.

Q10 noisefilters1

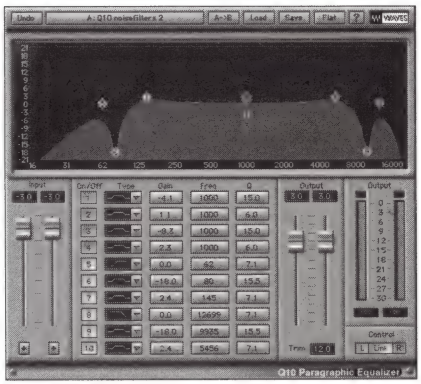

Q10 noisefilters2

The setups **Q10 noisefilters1** and **Q10 noisefilters2** provide a low-pass and a high-pass filter much better matched to human psychoacoustics. As may be seen from the filter curves on the Q10 graphic display, these filters embody several characteristics found to improve subjective tonal quality:

(i) they rise about 1.5 dB before the cut-off.
(ii) they fall to a deep notch.
(iii) they rise again to a reduced level of about -6dB outside the main band before finally falling away again.

This characteristic is very effective at reducing noise while still leaving enough of the high frequencies present (in the low-pass case) to give a much more natural sound quality than a simple roll-off. The two setups differ only in that the '1' filters have a shallower notch than the '2' filters.

The low-pass hiss noise filter uses bandgroup Q10 bands 8 to 10, and its frequency may be altered by selecting the three frequency value windows and dragging to either side as required, and the low-pass filter may be switched in or out by selecting the three In/Out buttons of bandgroup 8 to 10 and clicking on one of them.

The high-pass bass noise filter uses bandgroup Q10 bands 5 to 7, and its frequency may be altered by selecting the three frequency value windows and dragging to either side as required, and the low-pass filter may be switched in or out by selecting the three In/Out buttons of bandgroup 5 to 17 and clicking on one of them.

Additionally, two Superparametric dips are provided for additional selective noise filtering in **Q10 noisefilters1** and **Q10 noisefilters2**. One giving a 6 dB dip uses bandgroup Q10 bands 1 & 2, the other giving a 3 dB dip uses bandgroup Q10 bands 3 & 4. They may be tuned individually in frequency by selecting the respective two frequency value windows and dragging to either side, or they may be tuned together by selecting all four frequency value windows if both are used together to provide a 9 dB dip. The width of the Superparametric dip may be altered by selecting two or all four Q value windows and dragging to either side.

These effects are comb filters designed to notch out the fundamentals and first few harmonics of interference which occurs at a particular frequency.

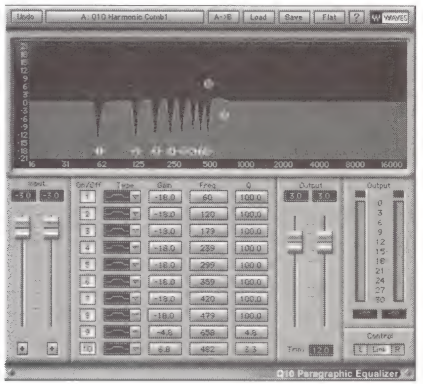

Q10 HarmonicComb1 notches out a fundamental frequency (seen in the top frequency value window) and the 2nd to eighth harmonic. Notch attenuation is around 15 to 18 dB.

Q10 HarmonicComb2 is the same as **Q10 HarmonicComb1** except that the notch depth is only around 9 dB, for use when less attenuation is required.

Q10 HarmonicComb5 is similar to **Q10 HarmonicComb1** except that only harmonics up to the fifth are removed, and the notch shape is more precisely controlled and narrower, in order to have less audible effect on wanted signals. The fundamental to fifth harmonic each occupy a bandgroup comprising 2 Q10 bands, respectively 1&2, 3&4, 5&6, 7&8, 9&10 for the fundamental and second to fifth harmonics. Individual notches may be switched out by switching out its bandgroup.

Q10 HarmonicCombOdd filters out the fundamental and odd harmonics only, from third to fifteenth harmonic.

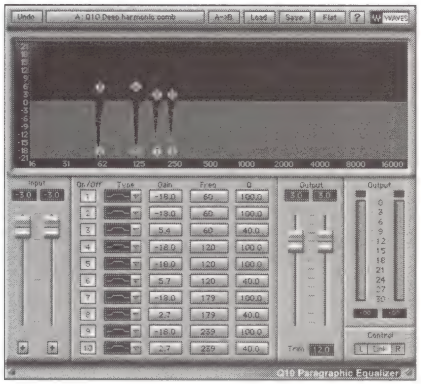

Q10 Deep harmonic comb

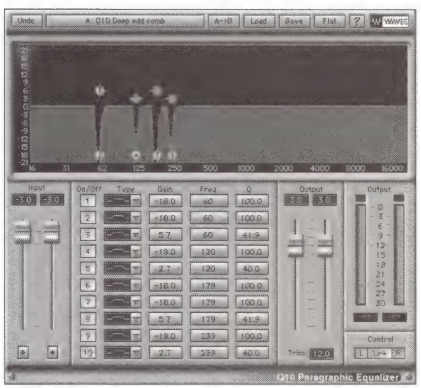

Q10 Deep odd comb

These effects are comb filters designed to notch out the fundamentals and first four harmonics of a particular frequency, such as a 60 Hz line power frequency interference or an acoustic "whine" due to high-speed motors.

The **Q10 Deep harmonic comb** setup notches out a fundamental frequency (seen in the top frequency value window) and the 2nd harmonic with a deep notch attenuation of around 30 dB, and the 3rd and 4th harmonics with a lesser notch attenuation in both cases of about 15 dB.

The **Q10 Deep odd comb** setup notches out a fundamental frequency (seen in the top frequency value window) and the 3rd harmonic with a deep notch attenuation of around 30 dB, and the 2nd and 4th harmonics with a lesser notch attenuation in both cases of about 15 dB.

In both cases, the fundamental frequency may be tuned up or down by selecting all 10 frequency value windows and dragging. When the setups are opened, the fundamental frequency is the US power frequency of 60 Hz, useful for attenuating recorded hum components. The setups may be modified to the European 50 Hz power frequency by European users by selecting all frequency value windows and dragging so that all frequency value windows are a multiple of 50 Hz.

These comb filters have some tonal effect on the sound, and are provided as an occasionally useful emergency "sound cleanup" method. They will not work when the pitched interference is a "buzz" with mainly very high harmonics.

It is common to do basic production work for multimedia 22 or 11 kHz sampling rates at the CD rate of 44.1 kHz and then to convert the sampling rate. Several tools for improving quality when doing this are provided. The following are bandlimiting (anti-aliasing) setups for files subsequently to be converted to 22 and 11 kHz sampling rates:

Q10 multimedia22
Q10 multimedia11

and these 3 files compensate for defects in the Sound Designer 2 sampling rate conversion process:

Q10 22.05postcomp
Q10 22.254postcomp
Q10 11.025postcomp

Note: These files are ONLY for use with Sound Designer II version 2.7 and below. Beginning with Sound Designer II version 2.8, the sample rate conversion for 22.050 and ll.025 are very flat and you should NOT use these post compensation setups.

Use of setups

Not for use with Sound Designer II 2.8 and greater. Only for SDII v2.7 and below.

These setups are intended to compensate for this droop by post-processing the sample-rate converted files at the final sampling rate.

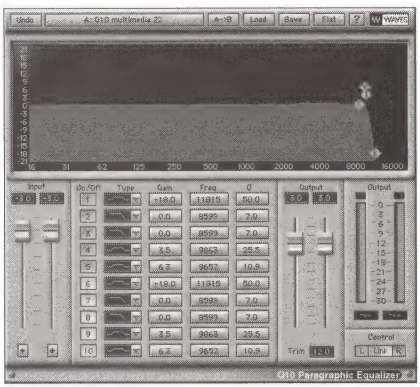

Q10 multimedia22

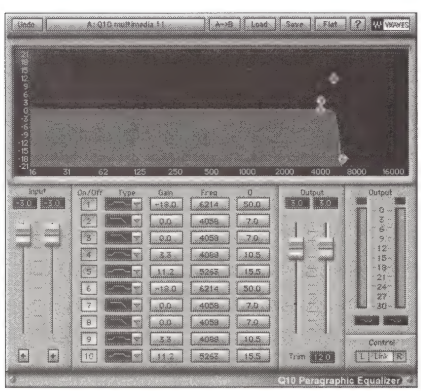

Q10 multimedia 11

The '22' setup file is simply a brickwall low-pass filter cutting off between 10 and 11 kHz. It is essentially flat up to 10 kHz, but falls away very rapidly above that frequency. The '11' setup is similar except that it is essentially flat up to 5 kHz and falls away rapidly at and above 5.5 kHz.

These setups have two applications:

(i) to monitor the sound of files at 44.1 or 48 kHz as they will sound when passed through the bandwidth of a 22 or 11 kHz multimedia sampling rate, and

(ii) to limit the bandwidth of a 44.1 or 48 kHz sound prior to subsequent processing for applications where the file will ultimately be converted to a 22 or 11 kHz sampling rate.

This is a desirable thing to do for two reasons:

(1) at all subsequent stages of processing, the sound will be heard -with the appropriate multimedia bandwidth, so that better production decisions can be taken.

(2) the bandlimiting will generally improve the sound of any subsequent sampling rate conversion used by removing high 'aliasing' frequencies that can degrade the sound.

As loaded, the brickwall filtering is provided by bandgroup Q10 bands 6 to 10. However, by switching an additional copy of the brickwall filter in bandgroup Q10 bands 1 to 5 in, the cut-off rate can be made extremely steep at the expense of only a small-degree of non-flatness of the in-band frequency response. In fact this cut-off is so steep that sampling rate conversion from 44.1 kHz to 22.05 kHz can be achieved after this filtering simply be discarding every second sample of the output signal file.

These setups provide:

(i) A brickwall filtering at 11 kHz (for the 22 kHz sampling rates) or 5.5 kHz (for the 11 kHz sampling rates) to minimize aliasing problems in sampling rate conversion.

(ii) A high frequency boost to compensate for the high-frequency roll-off in the Sound Designer II sampling rate conversion.

(iii) A gain reduction to prevent the significant high frequency boost from overloading and causing distortion.

You may monitor your 22 or 11 kHz kHz files more accurately by saving copies as a Macintosh SND Resource using the Save a Copy... command, and then playing these files back using your Macintosh's audio playback via an external high-quality monitoring system.

If using Sound Manager 3.0 or above, you can play the soundfile directly through the Mac system without any filtering from SDII. From the Setup menu, select Sound Playback. Check the Sound Manager box. Now if you click on the speaker icon in SDII, the selected region will play through the Sound Manager (and out the Macintosh speaker or sound output jacks).

It is hoped that a simple software processing tool to accomplish this will be available in the future to allow the Q10 to be used for processing 44.1 kHz soundfiles and converting them to multimedia files at 22.05 kHz.

SDII-SR Compensation

Q10 22.05postcomp
Q10 22.254postcomp
Q10 11.025postcomp

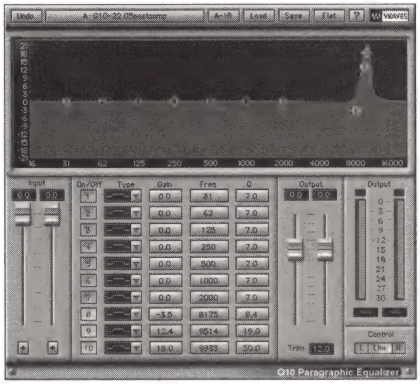

22.05postcomp

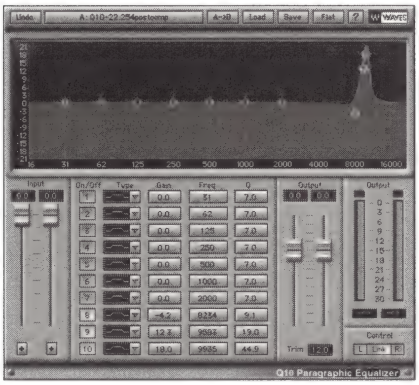

22.254postcomp

11.025postcomp

These setup files are intended to compensate for the poor frequency response of the sampling rate conversion algorithm provided with Sound Designer II **version 2.7 and below** when subsequently converting to the multimedia sampling rates of 22.05 kHz (PC) or 22.254 kHz (MAC) or 11.025 kHz from files sampled at 44.1 kHz. They typically increase the audio bandwidth at the final sampling rate by 25%, and with a generally flatter response up to cut-off.

Measurement in the digital domain shows that the Sound Designer II sampling rate conversion in **version 2.7 and below** from 44.1 kHz to 22.05 kHz causes a frequency response which is about 1 dB down at 7 kHz, about 3.5 dB down at 8 kHz, about 11 dB down at 9 kHz and around 30 dB down at 10 kHz. This drooping frequency response will cause a dull sound in multimedia applications. These frequency responses occur an octave lower when a further sampling rate conversion is made from 22.05 to 11.025 kHz sampling rates using the Sound Designer II sampling rate converter.

Chapter 7 - Distortion

The following setups produce intentional distortion effects which may be useful in certain creative applications, both in music production and in the creation of sound effects. They are generally more effective if the input sound file is first of all passed through a dynamic compression process.

Q10 E-guitar1

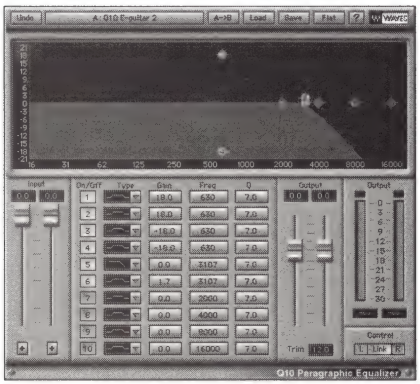

Q10 E-guitar2

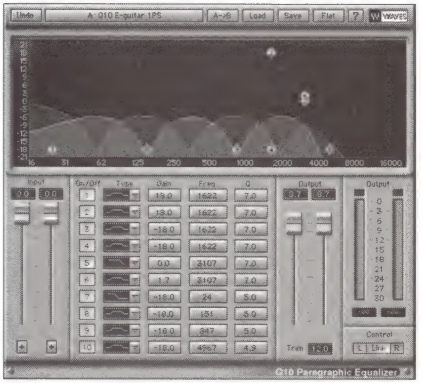

Q10 E-guitar1PS

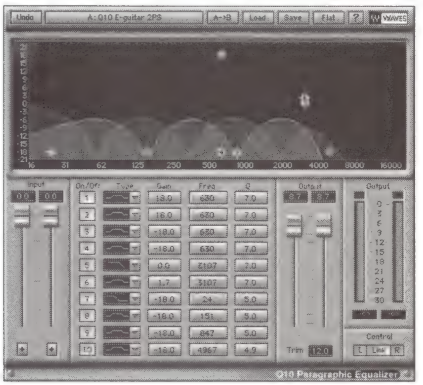

Q10 E-guitar2PS

These effects produce distortion of a kind that works particularly well with electric guitars, although there may also be creative applications when used with keyboards, bass lines, or even vocals! As loaded, the first E-guitar effect produces a distortion effect that retains a lot of the clarity of the original sound, whereas the second E-guitar effect is much more dirty and 'heavy metal/grunge'.

The E-guitar distortion effect is highly adjustable and tunable. The actual distortion itself is produced by bandgroup Q10 bands 1 to 4, and the high-frequency content of the distortion is tamed by a low-pass filter bandgroup Q10 bands 5 & 6.

The effects of particular adjustments are as follows:

The distortion quality is adjusted by selecting and dragging the frequency value windows of Q10 bands 1 to 4. This selects the frequencies in the original sound that are subject to distortion, and may be tuned to get the most desirable effect from a particular instrument or sound. With bass guitar, for example, generally a frequency in the low hundreds of Hz should be used, whereas electric guitar works best with from a few hundred to two or three thousand Hz, depending on the desired effect.

The filter may be separately tuned in frequency by selecting and dragging the frequency value windows of bandgroup Q10 bands 5 and 6, and adjusted for less or more 'toppy' distortion.

The distortion may be switched in and out without affecting the filtering by selecting and clicking the In/Out buttons of Q10 bands 1 to 4.

Q10 bands 7 to 10 are available for any additional EQ of the user's choice.

To reduce the amount of distortion without changing its character, reduce the input gains, and increase the output gains to obtain the required level.

The PS versions of the setups are identical except that a pseudostereo effect is added in bandgroup Q10 bands 7 to 10 to spread an input mono sound in the stereo stage. As with all pseudostereo effects, this will process a sound file with stereo output only if the input file is a stereo file with identical 'double-mono' channels.

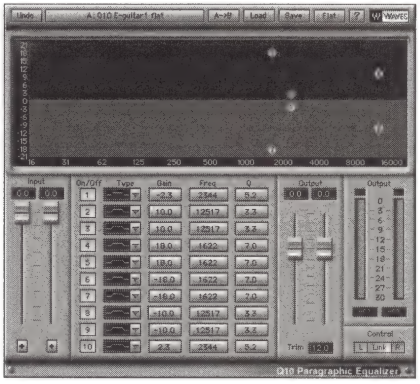

Q10 E-guitar1flat

Q10 E-guitar2flat

These effects are guitar distortion effects similar to the previous ones, except that the low-pass filtering is applied only to the distortion, leaving the original undistorted part of the guitar sound with a flat frequency response.

This effect uses all 10 Q10 bands, and the distortion-producing bandgroup is bands 4 to 7.

The distortion quality is adjusted by selecting and dragging the frequency value windows of Q10 bands 4 to 7. This selects the frequencies in the original sound that are subject to distortion, and may be tuned to get the most desirable effect from a particular instrument or sound. With bass guitar, for example, generally a frequency in the low hundreds should be used, whereas electric guitar works best with from a few hundred to two or three thousand Hz, depending on the desired effect.

To reduce the amount of distortion without changing its character, reduce the input gains, and increase the output gains to obtain the required level.

Chapter 8 - Tilt and Semi-Tilt

These setups, designed for 44.1 kHz sampling rate files, equalize sound files with a constant number of dB per octave boost or cut between 16 Hz and 20 kHz.

Bandgroup 1-5 gives a boost of the treble and cut in the bass. Bandgroup 6-10 gives a boost in the bass and cut in the treble.

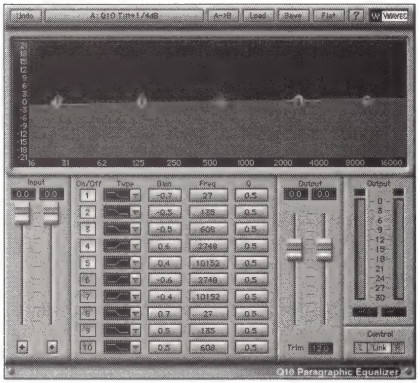

Q10 Tilt+1/4dB

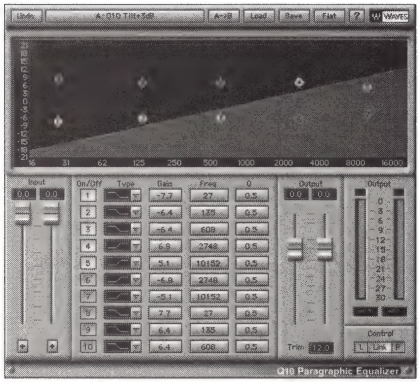

Q10 Tilt+3dB

The Q10 'Tilt' setup files give respective treble boosts of the indicated dB per octave if bandgroup 1-5 is switched in and treble cut the indicated dB per octave if bandgroup 6-10 is switched in.

The tilt setups are suitable for whenever an overall bass boost/treble cut or treble boost/bass cut is required without any particular frequency band being selected for special treatment. The overall effect is more "tonally neutral" than conventional treble/bass EQ, while giving an overall tonal tilt.

The **Q10 Tilt+3dB** EQ bandgroup 1-5 can be used to convert pink noise to white noise, or to convert white noise to blue noise and **Q10 Tilt+6dB** EQ bandgroup 1-5 can be used to convert pink noise to blue noise. Conversely, the **Q10 Tilt+3dB** EQ bandgroup 6-10 can be used to convert white noise to pink noise or pink noise to red noise, and the **Q10 Tilt+6dB** EQ bandgroup 6-10 can be used to convert white noise to red noise.

Bandgroup 1-5 of **Q10 Tilt+6dB** acts as an approximate differentiator within the audio band, and bandgroup 6-10 of **Q10 Tilt+6dB** acts as an approximate integrator within the audio band.

One can copy and paste two bandgroups from tilt setups to obtain other degrees of tilt. For example if one loads **Q10 Tilt+6dB**, and then copies bandgroup 1-5 of **Q10 Tilt+5dB** into bands 6-10 of **Q10 Tilt+6dB**, the overall effect of having bands 1-10 engaged will be to have a treble boost of 6+5 = 11 dB per octave. Obviously, especially with extreme degrees of boost, it may be advisable to reduce input gain until no clipping is audible, as the massive treble boosts involved will be likely otherwise to cause clipping.

In practice, these tilt setups also work well at a 48 kHz sampling rate, but not at lower sampling rates like 22 or 32 kHz.

Semi-tilt

Sub-classes of setups:

Bass-Tilt (tilt below a frequency 250 to 1 kHz),
Treble-Tilt (tilt above 2 kHz)
and **High-Tilt** (tilt above 4 kHz)

The Bass-Tilt setups, designed for all sampling rates, equalize sound files with a constant number of dB per octave boost or cut below a "hinge frequency" pre-set at 500 Hz, although the hinge frequency may be adjusted up to about 1 kHz.

Bandgroup 1-5 gives a cut in the bass.
Bandgroup 6-10 gives a boost in the bass.

Q10 Bass-Tilt 1.5dB

These setups, designed for all sampling rates, equalize sound files with a constant number of dB per octave boost or cut below a "hinge frequency" pre-set at 500 Hz, although the hinge frequency may be adjusted up to about 1 kHz. Bandgroup 1-5 gives a cut in the bass. Bandgroup 6-10 gives a boost in the bass.

The **Q10 Bass-Tilt** setup files give indicated bass cuts in dB per octave if bandgroup 1-5 is switched in and bass boosts in dB per octave if bandgroup 6-10 is switched in.

The hinge frequency below which the bass boost or cut occurs is loaded as 500 Hz, but this may be adjusted downward to around 250 Hz and upwards up to around a kHz or so by selecting and dragging the frequency windows of bandgroups 1-10. The hinge frequency can be set even higher than 1 kHz, but the characteristic departs from pure "tilt" at very low bass frequencies if one does so.

The bass tilt setups are suitable for whenever an overall bass boost or bass cut is required without any particular bass frequency band being selected for special treatment. The overall effect is more "tonally neutral" than conventional bass EQ, while giving an overall tonal tilt.

One can copy and paste two bandgroups from bass tilt setups to obtain other degrees of bass tilt. For example if one loads Bass-Tilt 4dB into Setup A, and then loads Bass-Tilt 3.5dB into setup B and then copies bandgroup 1-5 of it into bands 6-10 of setup A, the overall effect of having bands 1-10 of setup A engaged will be to have a bass cut of 4+3.5 = 7.5 dB per octave. Obviously, especially when boosting bass, it may be advisable to reduce input gain until no clipping is audible, as the boosts involved will be likely otherwise to cause clipping.

Treble and High-tilt

These Treble-Tilt and High-Tilt setups, designed for the 44.1 and 48 kHz sampling rates, equalize sound files with a constant number of dB per octave boost or cut above a "hinge frequency" pre-set at 2 kHz for the "Treble" setups and 4 kHz for the "High" setups. Bandgroup 1-3 gives a cut in the treble. Bandgroup 4-6 gives a boost in the treble.

The hinge frequency is not adjustable on these setups without altering the "tilt" shape.

Q10 Treble-Tilt 4dB

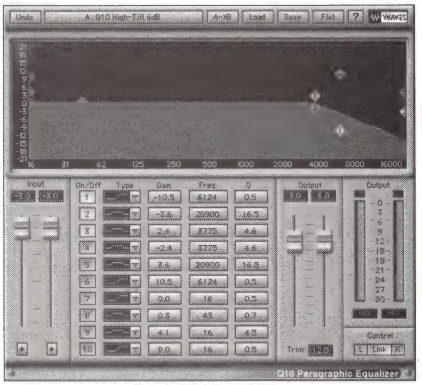

Q10 High-Tilt 6dB

The Q10 Treble-Tilt setup files give indicated treble cuts in dB per octave between 2kHz and 20 kHz if bandgroup 1-3 is switched in and treble boosts in dB per octave between 2kHz and 20 kHz if bandgroup 4-6 is switched in.

The Q10 High-Tilt setup files give indicated treble cuts in dB per octave between 4kHz and 20 kHz if bandgroup 1-3 is switched in and treble boosts in dB per octave between 4kHz and 20 kHz if bandgroup 4-6 is switched in.

The Treble- and High-tilt setups are suitable for whenever an overall treble boost or treble cut is required without any particular treble frequency band being selected for special treatment. The overall effect is more "tonally neutral" than conventional treble EQ, while giving an overall tonal tilt.

The High-Tilt setups are especially useful for subtle adjustment and enhancement or taming of high treble while having little effect on basic tonal quality, and are extremely useful for fine adjustment of lackluster or sizzly older recordings.

One can copy and paste two bandgroups from Treble- or High-tilt setups to obtain other degrees of tilt. For example if one loads **Q10 Treble-Tilt 4dB** into Setup A, and then loads **Q10 Treble-Tilt 3.5dB** into setup B and then copies bandgroup 1-3 of it into bands 4-6 of setup A, the overall effect of having bands 1-6 of setup A engaged will be to have a treble cut of 4 + 3.5 = 7.5 dB per octave above 2 kHz. Obviously, especially when boosting treble, it may be advisable to reduce input gain until no clipping is audible.

Combining Bass- and Treble- or High-tilts

Bass-Tilts may be combined with Treble- or High-Tilts. Load the Bass-Tilt setup desired in Setup A, selecting the bandgroup of 5 bands required and adjusting its hinge frequency as required, then load the Treble-Tilt or High-Tilt setup desired in Setup B. Select the bandgroup of 3 bands required; copy/paste it into an unused slot of 3 bands in Setup A. If Bass-Tilts with hinge frequency lowered to around 250 Hz are combined with High-Tilts, this can result in a tonally neutral effect combined with quite subtle alterations of bass and treble feel.

Chapter 9 - Pre/De-Emphasis

These setups are intended to add or remove various standard pre-emphasis curves from sound files in the digital domain. Since these are actually digital IIR filters, they are not quite exact simulations of the analog filters that define the curves, but they are usually very close.

It is important when using pre-emphasis (and sometimes de-emphasis) to reduce the input gains when necessary to prevent overload and clipping of the output.

Q10 CDpreemphasis

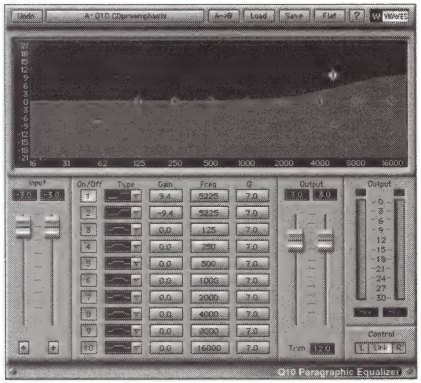

This setup approximates the standard CD and DAT pre-emphasis curve (with time constants 50 and 15 microseconds) at a sampling rate of 44.1 kHz, to within about 0.05 dB up to 16 kHz.

Band 1 alone implements the CD pre-emphasis on sound files recorded without pre-emphasis.

Band 2 alone implements the inverse de-emphasis, and may be used to remove pre-emphasis on a sound file having pre-emphasis.

Q10 48preemphasis

This setup approximates the standard CD and DAT pre-emphasis curve (with time constants 50 and 15 microseconds) at a sampling rate of 48 kHz, to within about 0.05 dB up to 18 kHz.

Band 1 alone implements the pre-emphasis on sound files recorded without pre-emphasis.

Band 2 alone implements the inverse de-emphasis, and may be used to remove pre-emphasis on a sound file having pre-emphasis. It is useful for removing pre-emphasis from tapes recorded on DAT machines using pre-emphasis.

Q10 FM-US

Q10 FM-US48

US FM pre-emphasis

This setup approximates the standard US FM broadcast pre-emphasis curve (with time constant 75 microseconds) at a sampling rate of 44.1 kHz (for Q10 FM-US) or 48 kHz (for Q10 FM-US48), to within about 0.1 dB up to 15 kHz.

The standard US FM pre-emphasis is also used for FM broadcasting in most of the Americas and in Japan.

Bandgroup 7 & 8 implements the FM pre-emphasis on sound files recorded without pre-emphasis.

Bandgroup 9 & 10 alone implements the inverse de-emphasis, and may be used to remove pre-emphasis on a sound file having pre-emphasis.

Bands 1-6 implement a brickwall low-pass filter cutting off at 15 kHz - the standard bandwidth limitation used in FM broadcasting.

A typical use would be to peak-limit sound files for FM broadcasting. To do this, first process a sound file by using on of two FM pre-emphasis setups, engaging bands 1-8, with a gain reduction on the inputs to avoid clipping. Then pass the resulting sound file through the Waves L1 limiter to limit levels of the pre-emphasized filtered signal. Then pass the resulting signal through bandgroup 9 & 10 ONLY of this setup of the Q10 to remove the pre-emphasis. The resulting file will then be limited in a way suited for FM broadcast.

Q10 FM-Euro
Q10 FM-Euro48

European FM pre-emphasis

This setup approximates the standard European FM broadcast pre-emphasis curve (with time constant 50 microseconds) at a sampling rate of 44.1 kHz (for Q10 FM-Euro) or 48 kHz (for Q10 FM-Euro48), to within about 0.1 dB up to 15 kHz.

Band 7 alone implements the FM pre-emphasis on sound files recorded without pre-emphasis.

Band 8 alone implements the inverse de-emphasis, and may be used to remove pre-emphasis on a sound file having pre-emphasis.

Bands 1-6 implement a brickwall low-pass filter cutting off at 15 kHz - the standard bandwidth limitation used in FM broadcasting.

A typical use would be to limit sound files for FM broadcasting. To do this, first process a sound file by using FM pre-emphasis engaging bands 1-7 of this setup, with a gain reduction on the inputs to avoid clipping. Then pass the resulting sound file through the Waves L1 limiter to limit levels of the pre-emphasized filtered signal. Then pass the resulting signal through band 8 ONLY of this setup of the Q10 to remove the pre-emphasis. The resulting file will then be limited in a way suited for FM broadcast.

Q10 RIAA
Q10 RIAA48

RIAA de-emphasis for phonograph records

This setup approximates the standard RIAA phongraph LP de-emphasis and pre-emphasis curves (with time constants 3180, 318 and 75 microseconds) at a sampling rate of 44.1 kHz (for Q10 RIAA) or 48 kHz (for Q10 RIAA48), to within better than 0.1 dB up to 15 kHz.

Its primary use is to equalize to flat soundfiles obtained by playing LP or 45 rpm phonograph record played on a velocity-sensitive cartridge into a pre-amp not having the RIAA equalization. It allows use of soundfiles from phonograph records without the need for an equalized pre-amplifier input. For the most accurate results, one should take care that the pre-amplifier input has an impedance matched to the load impedance recommended by the manufacturer of the phonograph cartridge used.

Bandgroup 1-2 is a rumble high-pass filter, set for -3dB point at about 20 Hz. This may be used to remove infrasonic rumble from soundfiles taken from phonograph records. Its operating frequency may be adjusted by selecting and dragging together the two frequency windows of this bandgroup. Use of this filter minimizes the risk of very high levels of infrasonic rumble once RIAA de-emphasis is used.

Bandgroup 7-10 alone implements the RIAA de-emphasis on sound files recorded with pre-emphasis, such as the output of a velocity phonograph cartridge feeding an unequalized pre-amplifier input. It may be used to equalize the sound from phonograph records to flat response. When this de-emphasis is used, take care to reduce the input gains, and also if necessary the output gains, to prevent output clipping

Bandgroup 3-6 alone implements the inverse RIAA pre-emphasis, and may be used to give RIAA pre-emphasis to a sound file having no pre-emphasis. There may not be much practical us for this, but it could be useful for the study of the effects of RIAA pre-emphasis.

One application is when older records such as 78 rpm records or early LPs have been played via a standard RIAA de-emphasis phongraph cartridge stage. You can engage the pre-emphasis Q10 bands 3-6 to remove the unwanted RIAA de-emphasis before adding a new de-emphasis for an older record de-emphasis characteristic.

Hint: For most older 78 RPM records pre-dating the late 1940s, a good approximation to older playback characteristics with constant velocity above a hinge frequency and constant amplitude at lower frequencies can be obtained as follows: Play the record over a conventional RIAA pre-amplifier stage and convert as usual to the digital domain. Engage Q10 bands 3-6 to remove the unwanted RIAA de-emphasis. In addition, engage Q10 bands 9 &10 to reintroduce constant amplitude characteristics at lower frequencies, BUT adjust the "hinge frequency" of this characteristic up and down by ear for best results by selecting the two frequency value windows of bandgroup 9 &10, and dragging these up or down. Q10 bandgroup 1 & 2 may be switched in or out to reduce low frequency rumble, and the frequency of the rumble filter may be adjusted as before by selecting and dragging the two frequency value windows of Q10 bandgroup 1 & 2. Q10 bandgroup 7 & 8 should remain switched out in this 78 rpm record re-equalization, or they may be used for additional conventional parametric equalization to the user's taste for additional tonal correction or for filtering "scratch noise".

Beyond what is said here, re-equalization of older 78 rpm records is a specialist skill, and there is no substitute for extensive experience for best results. The skills vary from finding the correct playback speed (older records were often recorded at other speeds), to the use of special stylus sizes and shapes (never use an LP stylus!) matched to the grooves of records from different times and places. With correct playback methods and skilled adjustment of equalization characteristics as above, one is often surprise to find just how very good old recordings can be–far better than is often realized when played back incorrectly. The natural sound quality of many old recordings can still put to shame many modern studio productions.

In order to avoid clipping of the Q10 equalization in this 78 rpm application, it may be advisable to reduce the input gain.

Q10 7.5 tape
Q10 15 tape

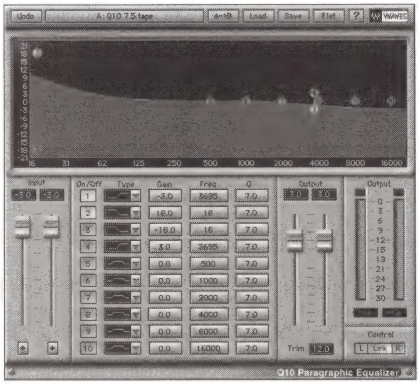

Conversion between IEC and NAB analog tape playback

Analog reel-to-reel tape at 7.5 or 15 inches per second speeds in the USA was largely recorded according to NAB playback equalization characteristics with time constants 3180 and 50 microseconds at both speeds, whereas in Europe, tapes were largely recorded for playback with IEC playback equalization characteristics with time constant 70 microseconds at 7.5 inches per second and 35 microseconds at 15 inches per second.

While many tape recorders offer playback with either NAB or IEC characteristics, other machines offer only NAB or only IEC playback, and will play tapes recorded with the other characteristic incorrectly. These setups provide correction of the equalization error caused by playing back tapes on the wrong equalization setting.

To correct European IEC tapes played back on NAB tape machines, select the setup for the tape speed (7.5 or 15). Select bands 1 and 2 only.

To correct US NAB tapes played back on IEC tape machines, select the setup for the tape speed (7.5 or 15). Select bands 3 and 4 only.

WARNING: These correction equalizations should only be used if no noise reduction is employed on the tape. Noise reduction will work correctly only if tapes are played on machines using the correct playback equalization.

These correction setups work at both 44.1 and 48 kHz sampling rates.

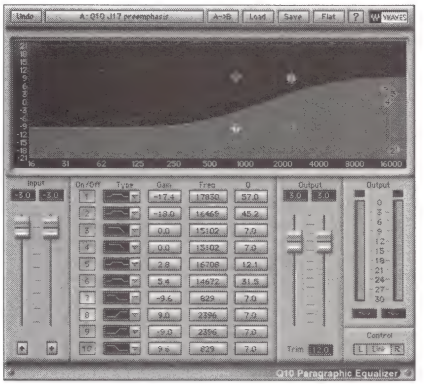

There are two internationally standardized pre-emphasis and de-emphasis characteristics for digital audio. The best known is the 50/15 microsecond characteristic sometimes used with CD and 48 kHz DAT tapes. However, there is another characteristic known by the technical name CCITT J17, widely used in broadcasting applications especially at a sampling rate of 32 kHz. J17 pre-emphasis incorporates a deep bass cut as well as a treble boost.

J17 has two time constants, 333.33 and 38.49 microseconds.

To add J17 pre-emphasis to a sound file without pre-emphasis, engage Q10 bandgroup 7 & 8 only.

To remove J17 pre-emphasis from a J17 pre-emphasized sound files, engage Q10 bandgroup 9 & 10 only. This implements J17 de-emphasis.

The J17 equalization will work well at sampling rates between 22 and 48 kHz, although its accuracy is optimized for the 44.1 kHz sampling rate.

This setup also includes a "brickwall" steep 15 kHz cut-off filter as bandgroup 1-6 to reduce the bandwidth to that used in broadcast NICAM applications. This filter is specifically optimized for the 44.1 kHz sampling rate, and SHOULD NOT be used at significantly lower sampling rates such as 32 or 22.05 kHz. The filter also still works reasonably well at the 48 kHz sampling rate.

On some sound files, clipping may occur, and if so, reduce the input gain. If the output is not fully modulated, some increase in gain may be desired–increase gain using the output faders to minimize risk of internal clipping in the Q10. The gain settings given are found to work well on many sound files, but may not always be best.

J17 is widely used at the 32 kHz sampling rate in broadcasting, and is the standard pre-emphasis used with the NICAM system of digital TV sound broadcasting in many parts of Europe, including the UK. Versions of NICAM are also widely used for digital distribution of broadcast sound signals both for TV and FM broadcasting, and has been used since 1972.

This setup can be used to convert conventional sound files to J17 pre-emphasis at a sampling rate of 32 kHz. Such sound files will be compatible with NICAM applications.

A typical use of this setup would be to limit sound files for broadcasting via NICAM or where NICAM links are used. To do this, first process a sound file by using NICAM pre-emphasis engaging bandgroup 7 & 8 of this setup, with a gain reduction on the inputs to avoid clipping if necessary. If the soundfile is at a sampling rate of 44.1 or 48 kHz, you may also wish to engage bandgroup 1-6 to remove frequencies above 15 kHz that are not wanted in the broadcast application. Then pass the resulting sound file through the Waves L1 limiter to limit levels of the pre-emphasized filtered signal. Then pass the resulting signal through bandgroup 9 & 10 of this setup of the Q10 to remove the pre-emphasis. The resulting file will then be limited in a way suited for NICAM broadcast.

Comments on other uses of J17 pre-emphasis

J17 is in theory internationally standardized, alongside the standard pre-emphasis used with CD and DAT, as a pre-emphasis for use with all digital systems, although it is rarely implemented except at a sampling rate of 32 kHz. Louis Fielder of Dolby Laboratories, noted authority on dynamic range issues, has argued that use of J17 pre-emphasis, unlike conventional CD pre-emphasis, has the effect of increasing the real dynamic range that can be handled by digital channels on real musical signals. The availability of the J17 emphasis (accurate to within about 0.1 dB up to 20 kHz at the 44.1 kHz sampling rate) will allow users to evaluate for themselves its performance, especially when the pre-emphasized files are quantized to low-bit wordlengths.

It is suggested that J17 may be ideal for handling multimedia files at sampling rates of 22.05 kHz, if played back via the corresponding de-emphasis. Indeed, J17 would make a good de-emphasis characteristic for multimedia files for lower noise and distortion. Fielder has also argued on the basis of extensive measurements of peak signal characteristics that J17 is also ideal for maximizing signal-to-noise ratio for professional applications at 44.1 and 48 kHz sampling rates, and would have made a far better pre-emphasis characteristic for CD than the one actually used.

The deep bass cut means that on many typical sound files having predominant energy below 1 kHz, the amplitude of the bass is greatly cut, and this allows the treble to be boosted considerably without the overall amplitude increasing. The inverse de-emphasis boosts bass and cuts treble, including any audible noise (including quantization noise) in the pre-emphasized signal channel. J17 is much better matched to many signals than other pre-emphasis characteristics.

Reference

L. Fielder "Dynamic Range Issues in the Modern Digital Audio Environment", Proceedings of the AES UK Conference "Managing the Bit Budget", (1994 May) pp. 3-19

Chapter 10 - Plateau

These setups are intended to provide detailed equalization of specific frequency bands without affecting other neighboring bands. They allow achievement of a "plateau" frequency response that is raised or lowered by a given number of dB within a user selected band relative to the rest of the audio spectrum. This plateau has steep sides (the transition between the gains takes about one fifth of an octave).

Plateau equalization is found particularly useful in bringing out or subduing particular instruments within a mix without too much effect on the tonal quality of the rest of the mix, so is particularly useful for subtle corrective rebalancing work. By allowing precise reduction or increase of particular frequency bands, it also allows one component of a mix to be reduced in spectral level in a critical frequency range to leave "space" for another component of a mix to be more audible, again without too much effect on overall sound of the mix.

Q10 Plateau+9dB
Q10 Plateau-3dB
Q10 Plateau+7.5dB
Q10 Plateau-4.5dB
Q10 Plateau+6dB
Q10 Plateau-6dB
Q10 Plateau+4.5dB
Q10 Plateau-7.5dB
Q10 Plateau+3dB
Q10 Plateau-9dB
Q10 Plateau+1.5dB
Q10 Plateau-18dB
Q10 Plateau-1.5dB

Q10 Plateau+3DB

The plateau equalizations are categorized by the relative level of the plateau relative to the rest of the audio band from +9 dB to -9 dB in 1.5 dB steps, plus a -18 dB plateau. Bandgroup Q10 bands 1 to 5 control the lower frequency step, and the frequency may be altered by select-and-dragging the frequency value windows of this bandgroup. The lower transition frequency is shown in the frequency value window of Q10 band 1. Bandgroup Q10 bands 6 to 10 control the higher frequency step, and the frequency may be altered by select-and-dragging the frequency value windows of this bandgroup. The higher transition frequency is shown in the frequency value window of Q10 band 10.

As loaded, these setups have transition frequencies 250 Hz and 4 kHz, giving a 4-octave wide plateau, but by altering one or both of these frequencies, the plateau may be varied in width from a fraction of an octave to many octaves. By this means, the plateau EQ can be applied to any desired frequency range.

Bandgroup Q10 bands 6 to 10 may be switched out if one just wishes for a single step transition up or down to the plateau.

One can even paste bands 1 to 5 from one Plateau setup with bands 6 to 10 from another Plateau setup in order to have different transition step sizes before and after the plateau. For example if band 1-5 are taken from the **Q10 Plateau+6dB** setup and bands 6-10 from the **Q10 Plateau+3dB** setup, then the first step will be from 0 to + 6 dB, and the second from +6 dB to +3 dB. Input and output gains should be adjusted to avoid overload when doing this.

By this means, the plateau setups allow arbitrary adjustment of two step transition frequencies and the step sizes (in 1.5 dB steps from -9 dB to +9 dB). This effectively allows equalization of the level of 3 frequency bands virtually independently of one another.

Octave setups

Q10 Octave+9dB
Q10 Octave-3dB
Q10 Octave+7.5dB
Q10 Octave-4.5dB
Q10 Octave+6dB
Q10 Octave-6dB
Q10 Octave+4.5dB
Q10 Octave-7.5dB
Q10 Octave+3dB
Q10 Octave-9dB
Q10 Octave+1.5dB
Q10 Octave-18dB
Q10 Octave-1.5dB

These are essentially identical to the Plateau EQ's except that, as loaded, these setups have transition frequencies 1251 Hz and 2502 kHz, giving an octave wide plateau.

Q10 Octave +3dB

This is a typical frequency range for picking out particular instruments, this particular octave being appropriate for electric guitars.

Plateau L/R types

Q10 PlateauLeft
Q10 PlateauRight
Q10 OctaveLeft
Q10 OctaveRight

Q10 PlateauLeft
Q10 PlateauRight (not pictured)

Q10 OctaveLeft
Q10 OctaveRight (not pictured)

Intended to be used with mono or double-mono signals, these four setups provide different plateau EQ on the left and the right channels. For the Left setups, the level of the left channel is boosted 3 dB and the right channel cut 18 dB in the plateau frequency region, The channels are the other way round for the Right setups. These setups in effect pan to the left (for Left setups) or right (for Right setups) over the plateau frequency region, and to the center at other frequencies.

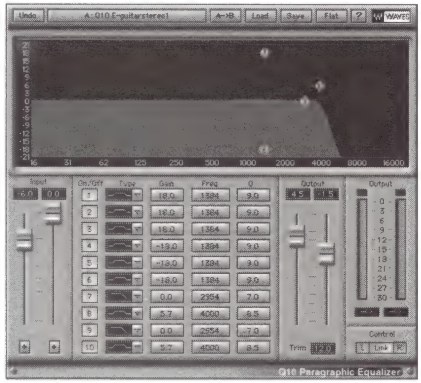

Q10 E-guitarstereo1

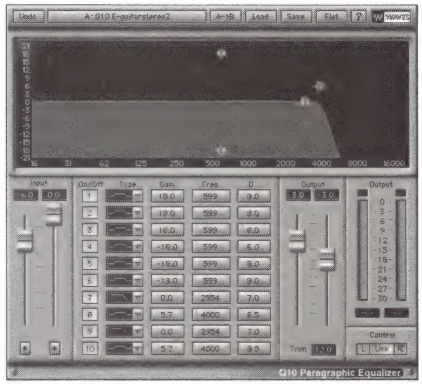

Q10 E-guitarstereo2

These are stereo distortion setups suitable for compressed electric guitar, and possibly bass guitar.

There are two parameters that can be tuned to alter the effect. The first is to select and drag (with linked stereo channels), the frequencies of bandgroup 1-6. This changes the quality of distortion. Higher frequencies, as in Q10 E-guitarstereo1, give a more "sizzly" effect, lower frequencies, as in Q10 E-guitarstereo2, a more "grungy" effect.

The second adjustment is to select and drag the frequencies of bandgroup 7-10. This adjusts the cut-off frequency in the treble of the guitar distortion effect.

One can also switch out bandgroup 9-10 in order to have less filtering of treble.

Chapter 12 - L/R setups

Q10 1.5octaveLR
Q10 2octaveLR
Q10 2.7octaveLR

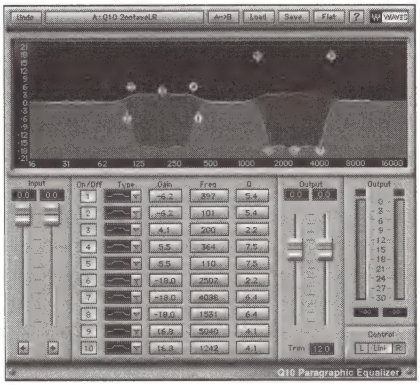

Q10 2octaveLR

These pseudostereo setups are designed to pan sharply defined frequency bands of sound to the left and to the right, with other sounds in the middle. They are designed for use with mono signals presented as double mono sound files. They differ from previous setups such as **Q10 PStunable**-types. These provide very sharply defined and relatively wide frequency bands in which sounds are panned to one side. The options are similar except that they differ in the width of the panned frequency bands, respectively 1.5, 2 and 2.7 octaves.

The bandgroup Q10 bands 1 to 5 controls the frequency band panned to the left. The bandgroup Q10 bands 6 to 10 controls the frequency band panned to the right. The panning of either bandgroup may be switched In/Out by selecting all 5 In/Out buttons and clicking. The frequency of operation of either bandgroup may be changed by select and dragging all 5 frequency value windows. The two bands may be in either order, i.e. bandgroup 1-5 may be at a lower frequency (as loaded) or at a higher frequency than bandgroup 6-10, but the two bands should not overlap too much–there is too much overlap if the left and right curve of the graphic display cross over much below 0 dB.

One may also have the two bands of different widths by pasting the bandgroup 6-10 of one of these setups to replace bandgroup 6-10 of another.

In use, this pseudostereo effect is particularly useful if the two bands are tuned so that the frequency ranges of particular instruments are panned to the two sides with the rest of sounds from the middle. This can often be used to synthesize a quite convincing stereo effect from mono source materials. Applications include for example historical mono material, mono PA feed signals from live recordings using PA and live ambient microphone signals, and mono ambience tracks in documentary and ENG recordings where ambiences may only have one track available. Mono compatibility should always be checked when this is important, since this effect is not always mono compatible.

Chapter 13 - Bass separation

Here are three demonstration files illustrating the effect of separation in the extreme bass.

Q10 bass-sepn Demo 1
Q10 bass-sepn Demo 2 (not pictured)
Q10 bass-sepn Demo 3 (not pictured)

It is commonly believed that the ears do not hear a stereo directional effect in the extreme bass, so that extreme bass can be fed to super-woofers placed almost anywhere in a room without degraded stereo effect. These demonstration files allow you to judge the truth or otherwise of this hypothesis for yourself!

All three files accept as input a dual-mono stereo file, and give as output a file in which the very low bass is panned to one stereo position and all other frequencies to another.

Demos 1 and 2 pan the bass to the left and other frequencies to the center. By switching bypass in and out, you can hear the effect of bass frequencies being removed from the center and panned to the left. The two demo setup files differ in the crossover frequencies used, which are respectively 75 Hz and 50 Hz approximately.

Demo 3 has crossover frequency of 80 Hz approximately, and places extreme bass at the center and other frequencies in the left.

All 3 demos give effectively constant power gain into the room at all frequencies.

The demos 1 & 2 use bandgroup 1-5 and demo 3 uses bandgroup 6-10, and the cross-over frequency may be adjusted by selecting the 5 frequencies in both channels in this bandgroup and dragging them up or down.

TrueVerb

Room Emulator and Reverb

Table of Contents

Chapter 1 ..... Quick Start 4

Definition of major terms 5
Quick Start on Thru mode 6
Example one 6
Example two 6

Chapter 2 ..... TrueVerb Interface 7

Time Response graph 7
Time Response controls 8
Frequency Response graph 9
Input/Output Section 10

Chapter 3 ..... Using TrueVerb .. 11

Thru mode 11
Using Thru mode 11
Adjusting Thru mode 12
Try this basic example 12
Send mode 13
Using Send mode 13
Adjusting Send mode 14

Chapter 4 ..... TrueVerb Tutorial .. 15

Some Reverb basics 15
How does TrueVerb work? 15
TrueVerb and stereo 16
Two modes of TrueVerb 16
TrueVerb in-depth example 16
Dimension 17
RoomSize 18
Distance 18
Link 19
Balance 19
DecayTime 20
PreDelay 20
Density 20
RevShelf 20
ERAbsorb 21
Reverb Damping 21
Evar and Rvar buttons 24

Chapter 5 ..... Multiple Parallel TrueVerbs .. 25
Multiple uses 25
Multi-Distances 25
Higher Quality 26
Combo Thru/Send 29
Surround Sound 29

Chapter 6 ..... Remastering Applications .. 30
Using TrueVerb for a live concert effect 30
Using TrueVerb for Remastering recordings 30
Natural EQ first 30
Proper stereo width 31
Natural sound+TrueVerb= maybe not 31
Using TrueVerb with previous Reverb effects 31

Chapter 7 ..... General TrueVerb Basics .. 32
How to begin 32
RoomSize and Distance 32
Imagine the room 33

Chapter 8 ..... Setup Library .. 34
About the setups 34
About the library 34
Thru setups 34
Send setups 34

Register Now!

We can not mail updates to you if you do not register.

Please make sure that you register your purchase by mailing your registration card or cards immediately, so that we can keep you up to date with our newsletter, new products and software upgrades. We welcome your suggestions and comments; email us at
support@waves.com.

Check our website —**www.waves.com** — for the latest versions of all our software. The web page also has demos of Waves applications and plug-ins, Frequently-Asked Questions and answers, User's Guides for all products in Adobe™ Acrobat™ (.pdf) format, and much more.

Chapter 1 • Quick start

You know you should read this manual very carefully, but you want to start immediately! So here's a brief two-page TrueVerb Quick Start. Following this is a description of the Thru and Send modes. After that, if you want a complete tutorial on TrueVerb, we suggest you read the cleverly-titled TrueVerb tutorial chapter. It offers detailed description of each control.

TrueVerb is very different from other reverb processors. It has two main modes:

[1] the 'Thru' mode, for True Distance Processing

[2] traditional 'Send' mode

You probably know how a send mode works, so we'll describe the 'thru' mode first: everything goes thru it. You don't 'mix' the sound of TrueVerb in with the original; instead, the original passes through it, so you put TrueVerb on the *Insert* point of a track, Aux, or Master.

Example One: you have a "board mix" of a concert, and it needs a 'room' feel. Patch the TrueVerb on the stereo *Insert* point of the Master fader. Determine how big the room is, how bright or dull it is, and, (here is the unique part) *determine how far the Distance to the stage is.* (This also works well for classical or jazz recordings that need a better room sound.)

Example Two: you have a submix of drums in your mix, and you need to put the drums in a 'good room'. *Insert TrueVerb* on your drum submix Master or Aux (this depends on your application); load the Drum Room (thru) setup file from the convenient TrueVerb Setup Library, and tweak. Instant room, without 'room mics' or some nice big hardwood-floor studio. (This works well for any submix, like vocals, guitars, etc.)

TrueVerb lets you create a true, natural-sounding room around a sound source, and control the RoomSize, frequency characteristics, and Distance. Or you can use it just like any other send-receive reverb and still use the early reflections and reverb together or separately. The Setup Library is made in two parts: Thru and Send. Put the TrueVerb where you want, choose the setup file you need, and tweak.

If you are unfamiliar with the Load/Save system, read the WaveSystem Manual for info on this and many more power tips common to all Waves plug-ins.

Definition of major terms

TrueVerb generates the virtual space by using two algorithms. The first is the Early Reflections, which generates multiple reflections in the first 100ms (one-tenth of a second) or so. The second is the Reverb, with many reflections, increasingly dense. By putting them together, a virtual room can be created. Some controls work on both algorithms, some separately.

Direct means the input signal, unchanged except in level.

EarlyRef means the early reflections, affected by the **ERAbsorb** control, the **Distance**, **DecayTime**, and the **Balance**.

Reverb means the output of the reverb processor, affected by the **RevShelf**, the **Reverb Damping**, **DecayTime**, and **Balance** controls.

Distance controls the perceived distance of the source.(This controls replaces the traditional "wet/dry" control).

DecayTime is the total 'reverb time' of the room you create, affecting the early reflections and reverberation.

RoomSize controls the size of the room in cubic meters.

Dimension controls the pattern of the Early Reflections to simulate a 1 to 4 dimensional space.

Balance controls the relative levels between Reverb and the Direct/EarlyRef, maintaining a constant perceived level.

Damping includes HF and LF (high/low frequencies) character of the Reverb itself. The HF and LF include adjustable Frequency controls (**HD Freq and LD Freq**) and **DampRatio** controls.

RevShelf is an input HF shelving filter for the Reverb.

ERAbsorb is an output HF filter for the Early Reflections.

The Freq control sets the frequencies for the RevShelf and the ERAbsorb.

PreDelay is the time before the Reverb begins, allowing the early reflections to 'reach' your ear first, much like in a true room.

Density controls the initial buildup of reflections in the Reverb.

This manual has nicely detailed explanations of all these points, and more. Please take the time to read it. Then you can really 'get the picture' on how to fully use the TrueVerb. For a more complete explanation on everything, read the TrueVerb Tutorial in chapter 4.

Quick Start on Thru mode

Example one

If you have a fully-mixed stereo file that is relatively 'dry', without reverb, then open the file on two tracks in your software. (A multitrack stereo mix will be fine too, without reverb). Ideally, you have some acoustic or jazz tracks, or even an 'off-the-board' mix from a concert.

For convenience to TDM users, we have provided a ProTools Session template in the Setup Library's Template folder called "QuickStart Example 1 Template". It already has the Master fader with TrueVerb inserted for you. Open this template if you wish, or follow the next 3 instructions. Of course, if you use the template, you'll have to import some stereo track to play with.
First, create a stereo Master fader.

Then, *insert* **TrueVerb** on that Master fader.

Last, click and hold the Load button on TrueVerb. A popup will appear; select the file "Medium Concert Hall ". Adjust the Distance control.

At first, this may sound very different from other reverbs you have used. You will hear the sound source moving away from you. TrueVerb is a mixing tool, so it simulates the distance while maintaining a constant perceived level. (In a real room, the level of everything would go down as you move further away.) You can't just move Distance around and pretend to be flying around in the room. Take your time and balance your monitor levels so that you can really 'get into' the space created by TrueVerb.

An excellent source for Medium Concert Hall is a nylon (classical) guitar, dry recording. By changing the Distance, RoomSize, and Absorption, you can define the room you want the guitarist to be in.

Example two

For convenience, we have provided a ProTools Session template in the Setup Library's Template folder called "QuickStart Example 2 Template". It already has the Aux fader with TrueVerb inserted for you. Open this template if you wish, or follow the instructions in the next paragraph. Of course, if you use the template, you'll have to import some stereo track to play with (preferably some percussion).

If you have a drum submixed to two channels, this is ideal, so mute all the other tracks and *insert* **TrueVerb** on the Stereo Master (if not already there). If the drums are on an Aux, then *insert* **TrueVerb** on that Stereo Aux of the submix. Make sure you route all tracks to the submix bus and not to the output; listen **only to the Stereo Aux output** and do not mix in any 'direct' sound from the individual drum channels to your Main Out.
Then, select the setup file "Drum Room " from the Load popup. You'll have a nice tight drum room that's not too bright or warm. Bongos and other solo percussion work famously well in this room too. It has high a Density and a short DecayTime.

Chapter 2 - TrueVerb Interface

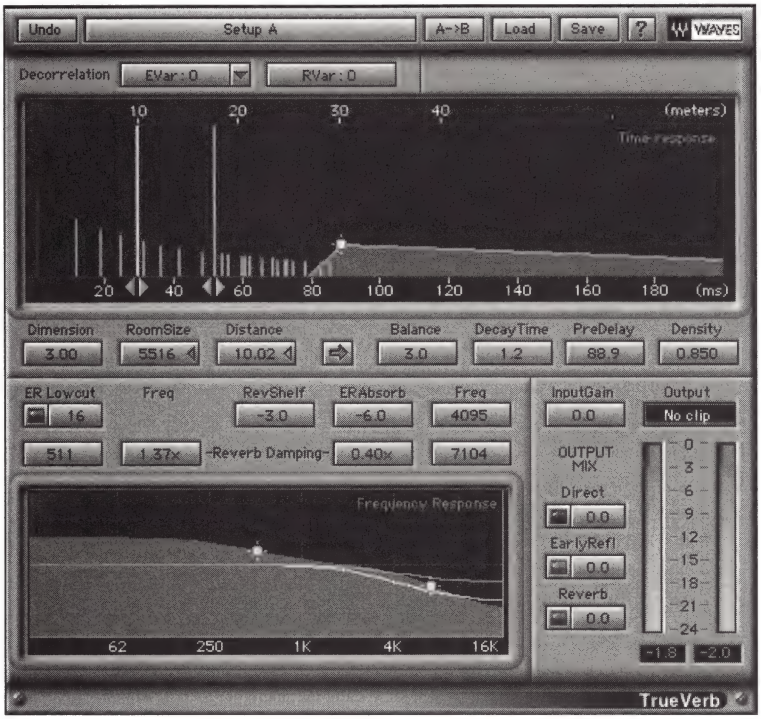

TrueVerb's graphic interface

Time Response graph

The top window (seen on the next page) is the Time Response graph, showing Room Size, Distance to source, the Early Reflections, PreDelay, Reverb tail, and all relative levels.

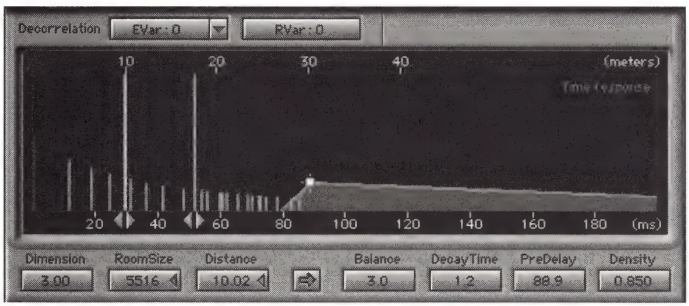

Room Size on the graph is shown at the top in meters (one axis) and the bottom in milliseconds; therefore, a size of 20 meters means a room 20x20x20=8000 cubic meters. The Room Size button, just below the Time window, shows the total cubic meter size. (Example: you want a small hall, 10 meters wide, 5 meter ceiling, and 20 meters deep = 1000 cubic meters).

Distance is also shown, in meters. Both Room Size and Distance can be adjusted by dragging the markers. You can even go outside the room, beyond the wall! In this case, the distance marker turns red to indicate 'out of bounds', just to let you know you're in the lobby.

The bottom edge of the Time Response graph is in milliseconds. From left to right, the graph shows the Direct sound (the first vertical bar on the left edge), the spread of early reflections on the left as white lines, then the Reverb, overlapping the reflections. As the PreDelay value changes, the Reverb shape will move to show the time offset.

Time Response controls

Just below the Time Response graph are the controls for the EarlyReflections and Reverb. When you click on a control, small arrows will show which way to move the mouse for adjustment, either left-right, or up-down.

Here are brief descriptions of each control:

Dimension Changes the 'character' of the Early Reflections to simulate an 'N' dimensional space; for example, a value of 2.0 will generate a pattern of early reflections that a 2 dimensional object might give; a value of 3.0 would correspond to a 3 dimensional room, and 4 vill be a virtual 4 dimensional space. Dimension does not affect the Reverb at all in tone, density, or diffusion.

Room Size Measured in cubic meters; for example, 8000 could be a room 20x20x20 (in meters), or 20x40x10. The blue triangles are a reminder that you can also change Room Size by dragging the marker in the Time Response graph.

Distance Measured in meters; controls the apparent distance from the listener to the input source. Automatically changes early reflections timing, and, if Link is on, also changes the reverb level and PreDelay. The yellow triangles are a reminder that you can also change Distance by dragging the marker in the Time Response graph. (When Link is off, Distance controls the Room level.)

Link Toggle button between Distance and Balance controls, which links the Reverb, Room, and PreDelay together. When on, the Reverb and PreDelay match the level of the last reflection in the Room. When off, PreDelay can be manually set, and the Reverb level can be changed ('unmatched' to Room).

Balance Controls the balance between the Reverb and the Direct+EarlyReflections levels. Positive values increase Reverb and decrease Direct+ EarlyReflections; negative does the opposite. When **Link is on**, the 0.0dB position represents the natural Balance between EarlyReflections and Reverb, given the rest of the controls settings. In this case changes of Balance will be arround this reference point. When **Link is off** the 0.0dB position results in equal levels for the EarlyReflections.

DecayTime Measured in seconds (standard -60dB reference); determines the rate at which the early reflections and reverb fade away.

PreDelay Measured in milliseconds (ms); controls the delay time before the reverb begins. When **Link is on**, the PreDelay is set internally, and no user adjustments can be made to it. When Link is off, you can freely adjust the PreDelay.

Density Ranges from 0.000 to 1.000; optimum smoothness at 0.850. Controls the density of the initial 'buildup' of reflections in the Reverb; higher values have higher density. Percussive sounds usually sound better with higher density values.

Frequency Response graph

The bottom window is the Frequency Response graph, showing the Reverb's frequency 'contour', the RevShelf, and the RoomAbsorb.

RevShelf Measured in dB; shown by the light blue line in the graph. A high frequency shelving filter that affects only the input to the Reverb section. The **Freq** sets the corner frequency for both the Shelf and the ERAbsorb.

ERAbsorb Measured in dB; EarlyReflections Absorption is shown by the yellow line in the graph. Affects only the Early Reflections, simulating the absorption characteristics of a real room's surfaces. For example, a dull room, with lots of fabric or carpeting, would have a very low value. The **Freq** sets the corner frequency of the absorption (for both the ERAbsorb and the RevShelf).

Freq Measured in Hz; sets the 'corner frequency' of both the Shelf and the ERAbsorb controls.

Reverb damping Actually a collection of four controls, which are unlabeled on the interface. The controls are listed below, from left to right on the interface. The frequency response of the Reverb's decay is controlled by them, and is shown on the graph as the blue area. Two small crosshair cursors (+) on the blue Reverb curve let you control both Freq and Ratio of either LF or HF.

LF Damp Measured in Hz; sets the corner frequency of the Low Frequency Damping Ratio (LF Ratio).

LF Ratio Controls the decay time of low frequencies in the Reverb section, relative to the **DecayTime** value. For example, a 'warm' room may have a value slightly above 1.00x; a room with better intelligibility usually has a value below 1.00x.

HF Ratio Controls the decay time of high frequencies in the Reverb section relative to the **DecayTime** value. For example, most halls might have settings between 0.25x and 0.50x. The higher the setting, the brighter the sound of the Reverb.

HF Damp Measured in Hz; sets the corner frequency of the High Frequency Damping Ratio (HF Ratio).

Input/Output Section

The bottom right corner of the window has all I/O controls and metering.

InputGain A full-range control, from 0.0 to silence (- infinity), adjusting the input level

Direct On/off toggle, with Direct level gain control

EarlyRefl On/off toggle, with Early Reflection level gain control

Reverb On/off toggle, with Reverb level gain control

VU meters TrueVerb's stereo output level is shown on the VU meters; below each meter is a numeric infinite peak hold, showing the highest value until reset. These peak hold values are sample-accurate. Above the meters is a clip indicator that will turn red and show the word 'clip' if a single sample is at full value (0.0dBFS). To reset the peak values, click directly on the meters, to reset the clip box, click directly on the clip box.

Chapter 3 - Using TrueVerb

Each platform uses TrueVerb in a different way, some by previewing and processing, some in real-time. The examples given here are specifically for real-time multitrack platforms. When used in preview/process applications, the Thru mode is generally assumed to be the method to use. Real-time usage will vary with desired effect, although most will use Thru. See your host application Notes and Read Me files for more information.

The distance effect of TrueVerb processing depends on the balance of these three components: **Direct**, **Early Reflections**, and **Reverb**. This balance is controlled by a single value on the TrueVerb (the Distance control). Of course, you can use TrueVerb in many ways, but to use the Distance effect properly, TrueVerb should be used in **thru** mode. The reason for this is simple: for the Distance effect to be most accurate, the level balance between Direct and Room signals is critical.

Thru mode

In this mode, the complete sound must pass through the TrueVerb. This can be either a complete mix, a submix, or an individual track. The output can be stereo or mono (but of course, the distance effect can not work in mono).

Inputs to TrueVerb can be stereo or mono. The output is a mix of **Direct** sound, **Early Ref** and **Reverb**. The Thru mode with stereo outputs is the mode to use for the most natural-sounding room acoustics.

Thru mode implies that the Direct output is switched in, plus the EarlyRef and/or Reverb. Normally you would use all three, but in cases where reverb already exists in the input signal you may wish to switch reverb out. In cases where you do not want a distance or room effect, but only to add reverberation, switch the **EarlyRef** out. Distance effect will work only with the **EarlyRef** switched in. For the distance effect to be most accurate, make sure the Gain controls for Direct, Early Ref, and Reverb are all set to 0.0.

Using Thru mode

The Setup Library is divided into two sections: Thru and Send. The setups make it easy to start using TrueVerb with any audio production.

As stated in the Quickstart section, the Thru mode means everything goes thru it; you should not mix the 'return' from the TrueVerb into the Output bus along with the original signal.

How do you make everything 'go thru it'? It all depends on your software and sound source. If the platform is a preview/process type, simply select the region or track you wish to process. If you use a Master, just insert the TrueVerb. If it is an Aux, route the sounds on a bus to the Aux, then Insert the TrueVerb on the Aux. If it is a single track, you may be able to Insert TrueVerb onto it, depending on your host software.

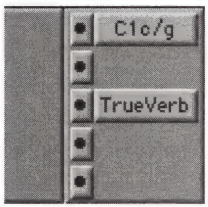

A TDM insert point with TrueVerb after the +C1.

Adjusting Thru mode

First, some important basics of TrueVerb Thru functionality:

- **Constant perceived level**
  Being a tool for multitrack mixing environment, one of the goals we had for TrueVerb was to maintain a (more or less) 'constant perceived level'. This is so that you don't have to work hard re-adjusting the mix each time you tweak TrueVerb. This is true for whatever changes you make to TrueVerb's controls.

- **Reverb Linked to Early Reflections**
  When Link is on, an arrow is shown on the button; the Reverb level and PreDelay match the EarlyReflections to actually let you directly create and control one integral acoustic space. In other words, when Link is on, you should think of TrueVerb as one integral algorithm simulating an acoustic space.

- **Distance & Balance, RoomSize & DecayTime :**
  These are the most basic TrueVerb controls. Each of them controls a very clear and intuitive aspect of your acoustic space, and each of them clearly affects all of the three TrueVerb components: Direct, EarlyRef and Reverb. After loading a setup, these are the first controls you should try out.

Try this basic example

We assume you have a stereo mix with a stereo Master. (For TDM users, there is a ProTools Session template in the Setup Library's Template folder called "Thru Template". If you use the template, skip the first step).

- Insert the TrueVerb on the Master Fader (or select it to affect the entire output).

- Mute all the tracks except for a single instrument, such as a guitar or voice, or the drums.

- Click and hold the Load button, and select the "Medium Concert Hall" setup in the top section of the popup menu.

- The Distance control is measured in meters. It will change the spacing of the early reflections, the PreDelay of the Reverb, and the levels of all three signals. You can see these changes in the Time Response graph as you move the Distance.

- Move the Distance control to a smaller value (closer). The Direct sound will be louder, the Room and Reverb will be softer. As Distance is increased, the balance of these three sounds will change. (The Distance control should not be mistaken as a simple wet/dry control, but it is similar). Please note that the overall sound level will not decrease as distance is increased (which it would do in reality), instead, as explained above, a 'constant perceived level' is maintained. Take your time to listen! It is really not possible to A/B and jump around with the Distance control (and honestly, quite hard to do in a real room)! It will take a few moments to adjust to the new reflections and levels.

- When Link is on and the Balance control is at 0.0dB, the level of the Reverb is matched to that of the EarlyRelections (the output controls for each section must be also set to 0.0 for this to be true). To increase the Reverb level, increase the Balance value (which will decrease the Direct+EarlyRef); to increase the Direct+EarlyRef level and decrease Reverb, decrease the Balance.

Experiment with other setups with solo instruments. A few good examples are: a drum kit in the "Drum Room", strings in the "Medium Concert Hall", vocals in the "Medium Plate", and rock guitar in the "Rock Guitar Room". The Setup Library has Rooms, Halls, Plates, and more.

Unmute the other tracks in your mix to put the entire stereo mix into the TrueVerb; select the Medium Concert Hall setup again to reset all settings. This will allow you to imitate a high-quality PA concert system, for example, to improve a "board mix" tape. Of course, some mixes won't work so well for this application. For detailed information on remastering in Thru mode, see the Remastering chapter in the Section 2-Applications part of the manual.

More traditionally, you may want to put only part of a mix into the TrueVerb. For example, you would route the desired sounds to a bus, then into an Aux channel with the TrueVerb inserted. An Auxillary is usually routed to the main output.

Here is an example: Audio tracks 1 and 2 have a premixed stereo soundtrack (to keep this example simple); Audio 3 and 4 have a stereo recording of a guitar solo and are routed to a stereo Bus (or Aux), panned hard left and right. The TrueVerb is inserted on the Aux. (Routing will vary with application).

If your application supports it, you could use several TrueVerbs on different tracks or submixes, each TrueVerb with a different room and distance setting. Each 'virtual space' will be overlaid into each other.

Send mode

Traditionally, reverbs are used by 'sending' signals across a bus to the reverb, then mixing the reverb output back into the main mix. Waves TrueVerb can certainly be used this way. Simply turn off the Direct sound (and maybe the Early reflections as well) and use TrueVerb in the way you might use other reverbs. Please note that the Distance effect will be compromised in this configuration; the critical level balance between Direct, EarlyReflections and Reverb is not under TrueVerb control. However, you can have different creative freedoms using the Room and Reverb outputs in a 'Send mode' configuration.

Using Send mode

The Setup Library has two sections: Thru and Send. Many of the setups are the same in the Send mode as their equivalents in the Thru setups. The differences in the setups are: Direct is switched out, Link is off, and the Reverb and Room levels may be different.

Adjusting Send mode

To give you the most flexibility, the TrueVerb is unlinked when you load a Send-type setup file. If you wish to keep the features of the Link mode, you can simply load a Thru-setup and toggle the Direct signal out.

For a basic example, we'll assume you have a stereo mix, with a Stereo Aux in place for the TrueVerb return. For TDM users, we have provided a ProTools Session file template in the Setup Library called "Send Template" (with a mono track 'send').

- If you're not using the template, or can't, insert the TrueVerb on the Stereo Aux (Effects Bus).
- Make sure to send the audio to the center (or feed two bus signals in TDM, one to the left, one to the right).
- Click and hold the Load button, and select the Medium Concert Hall from the top section of the popup menu.
- Turn off the Direct signal. To change levels of Reverb (and Room, if on): adjust the Balance.
- Adjust the Send and Return level to taste. Experiment with PreDelay and the RevTime for your mix.
- The EarlyRef can also be switched in and out; to change the character of the Room, adjust the Dimension, Distance, and RoomAbsrb controls.
- The other controls are fairly standard for Reverbs and are discussed fully in the chapter TrueVerb tutorial in the Section 2-Applications part of the manual.

Chapter 4 - TrueVerb Tutorial

To get the most out of TrueVerb, it helps to understand the basics of what a good reverb unit does. Before we go into more detail, it helps to understand basic terminology.

Some Reverb basics

A reverb unit provides at least three basic effects:

Life To make sound more "live" and "smooth" by adding sounds caused by reflections from the walls of a room. This makes the sound more like what you hear in "real life".

Space To increase the spaciousness of sounds. A good well-designed stereo reverb, like the TrueVerb, also adds a spatial effect making even mono sounds fill the space between stereo speakers. It adds a room effect that is has a wide stereoism. This allows mixes with little stereoism to have a wide and spacious effect.

Distance To place a sound in an acoustic space with a certain distance from the listener. Unlike most reverberation units on the market, the TrueVerb uses virtual reality techniques to place sounds convincingly within the actual space of a virtual room with good acoustics. You can adjust many characteristics of the room such as its size and reverberation time for almost any musical or dramatic requirement, and uniquely, you can also adjust where the sound is within the room by using TrueVerb's unique distance control.

How does TrueVerb work?

The effect of TrueVerb is achieved by a very precise combination of three sound components. These are Direct sound, Early Reflections, and Reverberation. The effect of an overall TrueVerb process depends on the balance and character of these three components. The TrueVerb's controls are linked together to simulate what happens in the acoustics of a good room.

Direct This is the original sound, arriving straight to the listener without any reflections from room boundaries. Because it is the sound arriving in a straight path from the original source, and so arrives first. The direct sound establishes the basic directional positioning of sounds within the stereo image. In the TrueVerb, there is no delay in the direct sound.

Early Reflections These sounds arrive at the listener after being reflected maybe once or twice from parts of the room, such as walls, ceilings and floor. They arrive later than the direct sound, often in a range from 5 to 100 milliseconds (thousandths of a second). The early reflections give your brain the information about the size of a room, and for the sense of distance of sounds in a room. They have an important role in determining the general character and sound of the room. The TrueVerb's Time Response graph shows the timing and level of the first 21 early reflections.

Reverberation After the early reflections, sounds reflect many times from parts of the room. There are so many they combine to form a continuous sound called reverberation, which decays slowly. This reverberation is subjectively different to early reflections, and adds a bloom and warmth to the sound, and also fills in gaps between sounds in a pleasant way. The time taken for the reverberation to die away is called DecayTime. In TrueVerb DecayTime marks the time by which Reverberation level falls 60dB. Decay time ranges from about 0.2 seconds for a small domestic room with plenty of carpets and soft furnishing to maybe 6 seconds for a grand cathedral.

TrueVerb and stereo

The TrueVerb process is a true stereo process at all stages. Every internal process within the TrueVerb is stereo-in/stereo-out, without any internal reduction to mono in any signal path.

This is much more than simply accepting stereo inputs and giving stereo outputs. All stages in between are also stereo, including early reflections and reverberation. This means that the TrueVerb process is different for each individual stereo position of sound in the input stereo mix, maintaining a distinct character for each stereo position. In this respect, it is closer to what happens in real rooms, where every sound in a different position in the room is treated a little differently by the room. Rooms do not add the same "blanket" effect onto every sound in the room; neither does TrueVerb.

If TrueVerb is fed with a mono input, it will produce a high degree of stereoism in the early reflections and reverberation, producing a sense of space and spread between the two stereo loudspeakers.

Two modes of TrueVerb

The TrueVerb permits two essentially different basic modes of use, with different advantages and disadvantages. These are Thru mode and Send mode.

They are discussed in **Section 1-Essential**, in the chapter **Using TrueVerb**.

TrueVerb in-depth example

In this lengthy section, we will describe each control in detail. Along with information about the control, we give some recommended settings, and discuss room acoustics related to that control, all at once. The goal is to provide information so that you —with good understanding— can design a room, then control the character of the room. The best thing to do is to keep a session open for experimenting while you read and work through the rest of this chapter.

For the following example, an ideal sound source would be a solo instrument, or a stereo mix of acoustic instruments without any other reverb. You should *insert* **TrueVerb** on the stereo Master Fader to the output. (You can load the Thru Template in the Setup Library's Template folder, if you wish, then import stereo tracks to play with.)

- Please reset TrueVerb by selecting the Medium Concert Hall from the Load popup.

- You can see in the Time Response graph all the components: Direct, Room and Reverb. The graph shows level (vertical scale) and time (horizontal scale). At the 0 delay time (far left) the red line represents the Direct sound and its level. Going from left to right, the 21 Early Reflections are shown by shorter white lines. Notice the decrease in level with increasing delay. The spacing is calibrated along the bottom in milliseconds and along the top in meters (sound travels approximately 1 meter in 3 milliseconds).

- To the right of the Room reflections is a blue-color triangular-area representing the Reverb envelope. It slightly overlaps the last of the Early Reflections. In reality, this is what a room sound is doing, with early reflections from the room surfaces becoming very dense, closely-spaced reflections known as reverb.

- Begin playback or preview. Turn the Early Reflections off by clicking its button in the lower right area and listen to the change. Turn it back on, then turn the Reverb off and listen to the change. Turn it back on before continuing to the next step.

About Dimension

Dimension affects only the Early Reflection spacing. It does not affect the Reverb in any way.

- Reset TrueVerb by selecting Medium Concert Hall from the Load menu again.

- Turn the Reverb off by clicking the Reverb button in the lower right area.

- A solo instrument track that has some dynamics is good for this (synthesizer pads would not). A percussion instrument or vocal would be good to work with.

- Change Dimension to 1.0. When Dimension is 1.0, you will hear a closed, and a slightly colored sound. In this case the early reflections have the simplest pattern because they are traveling in one dimension, like a wire or line, bouncing back and forth.

- Change Dimension to 2.0. Now the reflections travel in a two dimensional space, such as a plate. Note how the sound "opens up" as the reflection patterns becomes more complex.

- Change Dimension to 3.0. Reflections can now bounce around in a 3D space, like a normal room, you will hear a nice spacious sound, and see the patterns getting even more complex.

Dimension ranges all the way to 3.99, a virtual four dimensional space, with a very spacious sound and complex reflection patterns. Experiment with the value and use what works best for your application.

About RoomSize

- Move the RoomSize control to a shorter value. You'll see that everything in the Time Response graph slides to the left, linked together. The reflections changes relative levels, and timing between the reflections become shorter, which is what would happen in a real room. The Reverb also slides to the left (the PreDelay time becomes shorter) so that the Reverb always overlaps the last reflections.

RoomSize directly controls the perceived size of the room simulated by TrueVerb. RoomSize is measured on the button in cubic meters (1 cubic meter= about 35 cubic feet). It is also shown in the Time Response graph in meters of one dimension of the cube, and in milliseconds.

The Room Size is conveyed to the listener through the pattern of Early Reflections. As you increase Room Size you will hear the space opens up, for very small rooms (like a bathroom) a more colored sound will be produced. You can see the white lines —representing the Early Reflections' delay times and levels— change their positions and heights as Room Size changes. The PreDelay time will also change so that the Reverb stays overlapped with the last reflection.

Generally speaking, a room's decay time and size do not have to be related to each other. It is possible to have a small room with very reflective walls and a long decay time. Or, a large room with very absorbent walls and a short decay time.

Typical room sizes vary from about 50 cubic meters for a small room in a house, about 500 cubic meters for a small club, studio theater or intimate recital room, about 2000 cubic meters for a medium size theater or small concert venue, to 10,000 to 20,000 cubic meters for a full-size major concert hall.

About Distance

- Move the Distance control to a shorter value. You'll see that most of the energy in the Time Response graph shifts to the left. The Direct and first reflections become louder, the later reflections and Reverb becomes softer. Note how the Reverb level matches that of the last reflections, which is what would happen in a real room. TrueVerb actually matches its two different modules - Early Reflections and Reverb - to produce a single pattern of decaying reflections.

- Reset the TrueVerb by selecting the Medium Concert Hall again. Now change the DecayTime. You'll see the Reverb and the Room both change their characteristics so that the match between them is maintained.

About Link

The Link button controls whether the Reverb level and the PreDelay are automatically matched (overlapped) to the early reflections.

When it is on, the **Link** button shows an arrow indicating the link between the Early Reflections and Reverb controls. When Link is on, the **PreDelay** value is automatic, and the **Balance**, when at 0dB, will refer to a very smooth transition between Early Reflections and Reverb. Balance can still be adjusted above and below this 0dB point; PreDelay is automatic.

When Link is off, the Reverb level and PreDelay can be adjusted manually. In this case the Balance 0dB reference point is equal sound levels for the Early Reflections and Reverb.

- Turn Link off (the arrow will disappear from the button).
- To adjust the PreDelay, click and drag left/right on the PreDelay button. You'll see the Reverb envelope in the Time Response graph slide around to show the PreDelay setting.
- When unLinked, the Balance control gives you a very wide range of Reverb level adjustment. You'll notice that when unLinked, the Balance control will still affect the Room level.

Whether Linked or unLinked, the Balance control still balances between the EarlyRef and Reverb, while maintaining a constant perceived sound level. Note that the Link button is designed so that when you link and unlink, the actual reverb level does not change, rather the Balance control VALUE changes to give the same absolute gain under the new mode. Gain controls of Direct, EarlyRef, and Reverb must be set to 0.0 for accurate distance effect.

About Balance

Balance lets you control the relative levels of the Direct+EarlyReflections against the Reverberation tail. A nice feature of this control is that it will do this while maintaining the overall sound level of the TrueVerb.

Example:

- Reset the TrueVerb by selecting the Medium Concert Hall again. Now increase the Balance to 5.0. You'll see the Reverb level increase as the Direct and Room levels decrease. You may find that increased or decreased Balance values suits your taste! It's a good control to tweak first.

About DecayTime

DecayTime affects both the Reverb and the Early Reflections, controlling how fast they will die away. Different kinds of material require different DecayTimes. Speech is usually best with times of less than 1 second; popular music, jazz and chamber music are best with somewhat short times of about 1.2 seconds or less (with many exceptions). Orchestral music is best between 1.5 and 2.3 seconds, and choral or organ music is often best with longer DecayTimes of 3 to 5 seconds. Fast music generally works better with shorter DecayTimes than slow music. These are just a few points of reference for you to keep in mind.

About PreDelay

In the TrueVerb, as in most reverberation units, the early reflections and the later reverberation sound are produced by two different processes. The way these are matched together in level and time is important for the best effect. PreDelay changes automatically to keep the reverb matched to the last reflection. However, fully manual adjustment can be made if the **Link** control is off.

When unLinked, the PreDelay can be changed in the button numerically and by dragging the cursor on the corner of the Reverb envelope left and right.

TrueVerb measures predelay as the time between the direct sound and the start of the reverberation. So a predelay of 45ms means that the reverb starts 45ms after the direct sound. The shortest available predelay is 10ms.

Manually setting longer predelays than the Linked value generally gives an impressive "big room effect" but can also introduce an audible 'empty gap' between the EarlyReflections and Reverb that might upset musical timing and clarity. It's a non-natural effect, (but so is a lot of music these days) so the amount of predelay would be chosen according to creative needs.

About Density

Reverb is a collection of many reflections that become more and more dense. At the beginning of reverb, the initial density buildup can be varied with this control. Density ranges from 0.000 to 1.000. Higher values have higher density. Percussive sounds usually sound better with higher density values. A good general setting is 0.850.

About RevShelf

This is a tone control that affects the input to the Reverb section. It is a shelving filter, affecting frequencies above the value set in the **Freq** control, which is located to the right of the ERAbsorb control. In other words, everything above the Freq value can be turned down (a treble control!) before going into the Reverb. **RevShelf** is shown by the blue line on the Frequency Response graph.

The filter works only to decrease levels of these frequencies and is measured in dB. This control provides you with the option for a bright or dull reverb sound overall, separate from the sound of the Room. It is entirely possible to have a bright Early Reflections sound with a dull reverb, say, with a stage area with many hard reflective surfaces and an auditorium with lots of absorption. Or the other way around. It's your choice.

About ERAbsorb

As stated previously, the EarlyReflections and Reverb can have different frequency responses. ERAbsorb stands for Early Reflections Absorption. It is shown on the Frequency Response graph by the yellow line. ERAbsorb determines how the high-frequencies are absorbed along the EarlyReflections.

ERAbsorb is measured in dB. Its frequency is set by the **Freq** control, (the same control for the RevShelf). The high frequency absorption characteristics of parts of the room (walls, ceiling, floor) also affect the brightness and sense of space of a room due to their effects on the early reflections.

- Adjust the ERAbsorb value and listen to the effect on the early reflections. You might turn the Reverb off to make it more obvious.

- For best Distance effect, you should keep this frequency above 4000 or 5000 Hz. Lower absorption frequencies produce a duller sound which tend to remove information used by the brain to determine distance. The absorption shown in dB is the degree of absorption in the first of the early reflections. Later reflections, simulating several bounces of sound off the room boundaries, will have an increasing degree of absorption.

- Useful values are typically around -3 to -6dB absorption above 5 kHz. Subjectively, low absorption (that is, if set to 0 dB) gives a splashy spread of sound at high frequencies. This can sometimes be a desirable effect used to enhance a sense of liveliness of a room, but in other situations with high frequency percussion or bongos can result in distracting and unpleasant high frequency echo effects, in which case more absorption should be used. Otherwise, it's Bongo HF Fury, so to speak.

The controls on the TrueVerb allow the frequencies of the HF absorption of the Early reflections and of the HF damping of the reverb to be linked and adjusted together. To do this, shift-click on each of the controls, RoomAbsrb and HF Damp, both will respond to the mouse or arrow keys. (In fact, this is true about all controls on Waves plug-ins with most applications. Read the WaveSystem Manual for complete details).

About Reverb Damping

Damping is a group of 4 controls and is part of the most essential adjustments to realistic-sounding reverbs. The controls are grouped together in a line, immediately above the Frequency Response graph. They are unlabeled.

So, from left to right on the TrueVerb interface, you'll see the low frequency value (Hz), the LF ratio, then the HF ratio, and finally the high frequency value (Hz). This arrangement allows for easy modification of both Ratios at the same time, by shift-clicking them and dragging to change values.

You can also change damping by dragging the crosshair cursors that "ride" on top of the blue Frequency Response shape. Each cursor lets you control the frequency and ratio simultaneously.

Damping is basically the frequency response of a room's reverberation. The total DecayTime of a room might be 2 seconds, but high frequencies might have a DecayTime of 1 second. In this case, the high frequencies would have one-half the total time. This can be expressed as a Ratio. For this example, the HF ratio would be 0.5 (so, 2 seconds multiplied by 0.5 = 1 second). A room with hard, reflective walls, such as a bathroom with tiles, can have a HF ratio greater than 1, giving longer decay times (and a very bright sound).

HF damping is the reduction of reverberation time at high frequencies, due both to absorption of high frequencies by air and absorption in the walls. Generally speaking, this high frequency reduction in reverberation time is necessary to avoid a very splashy or sibilant sound. The degree of HF damping actually changes in real rooms according to furnishings, audience size, air temperature and humidity.

- Typical HF damping frequencies may be between 4 and 7 kHz.

- A typical damping ratio may be around 0.25x to 0.5x. Very dry air has greater HF damping, and stone walls, such as in many churches, give less HF damping than do wooden or furnished rooms.

LF damping is the reduction of reverberation time at low frequencies, due to low frequency transmission in the walls. Most rooms with solid walls have little low frequency damping. In general, this is because a wall with more LF transmission has a greater level of "noise nuisance" caused by bass sounds leaking out of the room and outside bass sounds leaking into the room. If noise nuisance isn't a problem, (and it won't be in a virtual room!) then LF damping increases the clarity of sounds, giving a clearer bass, especially on speech and orchestral music. Therefore in the TrueVerb, the ability to adjust LF damping for optimal clarity is provided.

- Start with a damping frequency around 300 Hz and an LF damping of about 0.7x. Too much LF damping sounds unnaturally thin and loses warmth.

- Many rooms actually have a longer reverberation time in the bass than at middle frequencies, and if you want some extra warmth, set the LF damping at slightly greater than 1.00x.

For even more detailed suggestions about damping (!), read the General TrueVerb Tips, Chapter 8.

The low and high damping each have two controls: a cutoff frequency value and a ratio value. The ratio is a multiplier of the DecayTime. The ratio is shown from 0.10x to 2.00x. For the low frequencies, the decay time at lowest frequency (DC) will be the LFRatio multiplied by the DecayTime; at the cutoff frequency the decay time will be half way in between, returning gradually to DecayTime above the cutoff frequency. The high frequency damping has the same type of controls, except that all the frequencies above the value are affected.

Example:
If high frequency damping is set to 5000, its ratio is set to 0.4x, and DecayTime is set to 2.0, then the frequencies above 5000Hz will have a decay time of about 0.8 seconds in the Reverb.

HF damping is the reduction of reverberation time at high frequencies, due both to absorption of high frequencies by air and absorption in the walls. Generally speakiö◊, this high frequency reduction in reverberation time is necessary to avoid a very splashy or sibilant sound. The degree of HF damping actually changes in real rooms according to furnishings, audience size, air temperature and humidity.

- Typical HF damping frequencies may be between 4 and 7 kHz.

- A typical damping ratio may be around 0.25x to 0.5x. Very dry air has greater HF damping, and stone walls, such as in many churches, give less HF damping than do wooden or furnished rooms.

About EVar and RVar buttons

In order to have the best quality with multiple TrueVerbs (for platforms that support multiple parallel inserts with the same source), there are two controls that provide small variations in the Early Reflections and the Reverb. The variations are in the reflection patterns of each, with slightly different audible characteristics. They are primarily for multiple-TrueVerb uses, but you can use these controls anytime you wish. Their effects are subtle.

Each control defaults to a value of 0. To change to a variation, click one of the buttons.

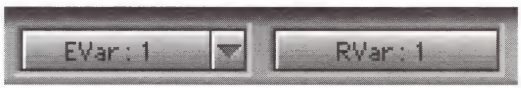

The **Rvar** button has 2 values, 0 and 1 (Standard + 1 variation). The **Evar** button has 7 values, 0 through 6 (Standard + 6 variations). You can listen to just the EarlyRef by switching out the Direct and Reverb, and hear the subtle changes in the early reflection patterns.

For small shifts in the character of a room, you can change these at anytime during mose uses, unless you are using multiple TrueVerbs, in which case you should be careful to follow the recommendations in the **Multiple Parallel TrueVerbs** chapter.

Chapter 5 - Multiple Parallel TrueVerbs

These variation values are saved with any setups you make. Some setups in the Library use different EarlyRef patterns.

An advantage of TrueVerb is that on some platforms you can run multiple copies of the processing (limited by the processing power available).

The ability to run several copies lets you do things that would be quite expensive to do with conventional reverb units. In this chapter we describe some particular uses of multiple copies.

Multiple uses

Some applications include:

Multi-distance Different sounds or submixes can be placed at different distances within the same virtual room.

Higher quality An increased-resolution reverb can be created by using the extra processing power from an additional TrueVerb.

Combo A simultaneous use of Thru processing for optimum distance effect plus Send/Receive mode for extended control of reverb.

Multi-Distances

If you want some sounds to be closer than others within the same acoustical space, the solution is to use two (or more) copies of TrueVerb. This example is with two TrueVerbs, for two independent distances.
The idea is simple. Prepare two submixes which we will call A and B. Some sounds are mixed into stereo submix A and others into stereo submix B, but in general you should not assign any given track to both submixes - only to one or the other. The final output mix will have the two individual submixes A and B with added TrueVerb processing.

These examples of multiple-parallel TrueVerbs are easist to demonstrate within TDM. So, we have provided a ProTools-TDM template in the TrueVerb Setups Library, Template folder, called "Multi-Distance Template". You can load this for a quick example.

Here are the details of that template. We have assigned all tracks for stereo submix A to bus 1-2 and all tracks for stereo submix B to bus 3-4. Two stereo Auxs have Bus 1-2 and Bus 3-4 as their inputs, respectively. Each stereo Aux has a TrueVerb inserted. The stereo Auxs outputs are assigned to Outputs1-2. (This makes Busses1+3=output 1, and Busses2+4=output 2).

Mixing the session is the same as usual, except that you'll assign the output of tracks to either Bus 1-2 or Bus 3-4.

Using two submixes has many advantages, allowing different components of the final mix to be rebalanced and altered without having to redo the whole mix again. For example, if at a future time you wish to do a surround-sound mix, you can go back to the two stereo submixes and place each of them across different speakers in the surround mix. Stereo submixes can be saved as multiple tracks on a digital multitrack such as an ADAT™ or DA-88™.

So, in other words, we now have this setup: Stereo Aux 1 has submix A, with a TrueVerb, and Stereo Aux 2 has submix B, with a TrueVerb.

It is possible to put submix A and submix B into two completely different acoustic spaces, but this might be very confusing to the ear and brain - certainly not at all natural. There's no accounting for taste or the creative urge. Do you what you wish! Just be aware that the most natural effect will be obtained if the two TrueVerbs have the same settings (the same acoustical space) but simply have different distances settings. Then you will have the effect of all sounds being in the same room but in different locations.

Essential steps for multi-distance:

- Choose a setup file to work with, or create your own on Submix A's TrueVerb. Select the distance you want for the sounds on submix A.

- Copy the submix A TrueVerb settings by selecting the TDM 'Copy Settings' command from the Effect menu (shortcut is shift-command-C).

- Bring the submix B TrueVerb to the front. Paste the settings by selecting the TDM 'Paste Settings' command from the Effect menu (shortcut is shift-command-V).

- Adjust the **Distance** in each TrueVerb to give the desired distance for the submixes.

Obviously, (well, we hope it's obvious), you can use three or even more of these TrueVerb submixes for realistic placement in space of single instruments, stereo images, subgroups, or other sources, in a single room setup (or different rooms). Be creative. It's your job!

Higher Quality

One of the limitations in the sound quality obtainable from the TrueVerb is because of limitations in the amount of processing power available. Any true-stereo process is confined to a single DSP chip (for non-Native apps) at the time of this writing. While careful design has maximized the quality within this limitation, the use of two TrueVerbs in parallel can give improved quality for TDM users.
The idea is to take a stereo mix and feed it to two copies of TrueVerb running in parallel, for example, on two stereo Auxs with the same input and output busses. This is very similar to the multi-distance application, except that there are not two separate submixes.

There is a Pro Tools session template for this example in the Setup Library, Templates folder, called "Higher Quality template". This simple template assumes you can route a stereo mix to Bus 1-2 (instead of Outputs 1-2) by simply importing an already mixed track, or you can add more tracks and import individual tracks for a regular mix. (In this case, it would be easier to just modify a copy of your session, of course.)

- If you don't use the template:
  create two stereo Auxs with inputs set to bus1-2;
  route all tracks to bus1-2 instead of Out1-2;
  route both Auxs outputs to output1-2;
  insert TrueVerb on each stereo Aux.

- Next, **mute** the Aux with TrueVerb-B.

- Switch to TrueVerb-A, and load a setup to start with, if you wish.

There are two procedures for getting improved reverb from this template.

Procedure One

In this case, we intend only to improve the sound quality of the **Reverb** output. To do this, we will use only the Reverb variations (**Rvar**).

- Make sure that TrueVerb-B is muted.

- Adjust the settings of TrueVerb-A until you have a satisfactory effect, including Distance, Reverb, and so forth.

- Then, decrease the **Balance** value by 3 dB.

- Copy the setup of TrueVerb-A (shift-command-C, or select Copy Settings from the Effect menu).

- Paste the setup into TrueVerb-B (shift-command-V, or select Paste Settings from the Effect menu).

- Turn off the **Direct** and the **EarlyRef** in TrueVerb-B. (Only **Reverb** will be on).

- **Now:** hold the Command key down, then click the **Rvar** button in the upper left corner of TrueVerb-B. It will show a value of '1'.

- Unmute the second Aux channel. It is best to group the two Aux channels for easy level adjustments (they are already grouped in the template.)

This will have exactly the same reverb effect you set up on TrueVerb-A, but now with improved quality of reverberation due to the second TrueVerb.

Procedure Two

Use this method when you wish to have both a higher density of early reflections and an improved quality of reverb.

- First, find a satisfactory reverb effect on TrueVerb-A (with copy B muted).

- Copy the setup of TrueVerb-A and paste into TrueVerb-B. To do so, use the copy/paste settings commands from the Effect menu of ProTools.

- Turn the **Direct** output off in TrueVerb-B.

- In TrueVerb-B only, click each of the **Variation** buttons in the upper left corner (**Rvar** and **EVar**). Each button will be orange and show a value of '1'. If you wish, you can command-click the Evar button repeatedly and choose any of the 6 variations.

- Unmute the second Aux. Use the same gain on both Aux channels. (Hint: group the Aux faders. They are already grouped in the template.)

- This arrangement has twice the number of early reflections (a total of 42). It also has the effect of changing the effective Room Size, decreasing it by a factor of 2. Therefore, if you had a RoomSize of 8000, using two TrueVerbs this way gives you a real RoomSize of 4000.

- Please note that if you change any settings in TrueVerb-A again, you must copy-and-paste the settings over to TrueVerb copy B, and then set both of the Variations back to '1' in TrueVerb-B. Don't forget to do this! The Quality Reverb Commission will knock on your door, or something like it. Quality is of prime concern in this arrangement.

The overall result is not only a smoother reverberation tail, but a more dense set of early reflections than in either the standard single-copy TrueVerb or than when using Procedure One. The disadvantage compared to Procedure One is that it is a bit more complicated to put together.

Combo Thru/Send

As you should know, TrueVerb has two modes of use, Thru and Send. Thru mode is the one to use for achieving the most accurate room simulation effect. Send mode is best where one simply wishes to have a high degree of flexibility in the amount of reverb on each different component of a mix.

However, by using two copies of TrueVerb, it is possible to combine many of the advantages of the two modes, although setting it up becomes more complex.

Here's the arrangement: again, two TrueVerbs at the same time, but this time, the first in Thru mode, the second in Send mode (Direct and EarlyRef are out). The first (Thru) is inserted on a Master Fader (or wherever the stereo mix or submix is); the second (Send) is on a stereo Aux to receive a send mix from any channel in the session.

The TrueVerb-Thru produces the distance simulation using Direct and EarlyRef, but adds little or no reverberation; TrueVerb-Send has Direct and EarlyRef turned off, and provides different amounts of reverberation for different sounds in the mix. Simply send the desired signal across the bus to the TrueVerb-B.

We have created a simple ProTools session template in the Setup Library Templates folder called "Combo template".

Surround Sound

Although TrueVerb is a true-stereo device, it is possible to use two copies of it to produce a true surround reverb for surround-sound mixes. This can be useful when doing mixes for multichannel sound both for music and post-production work.

In principle the idea is simple. Suppose that you are mixing a standard stereo image (in front of you), but wish to add a full surround room effect. The idea is to use a TrueVerb to produce a reverb to be fed to the front stereo stage, and a second TrueVerb to produce a separate distinct reverb sound to be fed to the surround loudspeakers.

Of course you could feed two completely different reverb effects to the two sets of loudspeakers, and get some kind of surround effect. But by careful matching of the two copies of TrueVerb, it is possible to retain the accurate room simulation properties of TrueVerb, placing the sound within a complete surround room presentation.

How to use for Surround sound:

The setup is exactly the same as Procedure Two for the Higher Quality arrangement in the previous section. The only difference is that the second TrueVerb is routed to rear speakers (say, to Outputs 3-4). The Direct is switched out on TrueVerb-B.

Chapter 6 - Remastering Applications

Using TrueVerb for a live concert effect

One particularly useful application for TrueVerb (in Thru mode) is in reprocessing PA-desk ("board mix") live recordings to sound like actual performances — but with good quality PA system.

The common problem is the only recording of a good live performance often is a recording of the PA mixing board, which was never meant to be heard except through the system, and of course does not have the atmosphere and "sound" of a live event.

Often this recording includes reverb effects from the PA board and rack itself, which may not be the most subtle. (Hey, it's a concert.) There is nothing much you can do to remove these effects, but what you can try and do is to turn the sound into one that might have come from a very high quality PA system.

The TrueVerb provides a solution to making such recordings more acceptable.

You simply have to dial up the actual Room Size of the concert location, reverb characteristics and the desired distance. You will have to do some EQ to the dry signal first to get the most natural sound balance possible. After that, you can place it in the big concert hall with a super-quality-PA-system-sound by using the TrueVerb.

Using TrueVerb for Remastering recordings

More generally, the TrueVerb can also be used to remaster demo and poor studio recordings, by giving them a more natural acoustic environment. This is not something that conventional reverbs can do well. The idea is simply to try and choose a distance, room size and reverberation characteristic that are compatible with the music, but which do not create an obvious "added-reverb" effect.

Ideally the effect will seem quite subtle on direct A/B comparison.

It is a characteristic of human hearing of the effect of acoustical perspective and distance that the effect is most obvious after listening for a longer time, often taking at least several seconds to establish itself after A/B switching. A/B switching is not a good way of determining the effect on prolonged listening of TrueVerb Thru processing.

The effect of good TrueVerb thru-processing of a recording which initially sounds flat and crude is that you will have lower listening fatigue on prolonged listening, and the sound components seem to blend together in a more natural way. This is not the typical effect of conventional reverb, which does not provide the ears with the correct spatial cues.

For best results with TrueVerb Thru-processing, the following points are well worth remembering:

Natural EQ first

Try and EQ the recording to have the most "natural" balance possible before adding TrueVerb processing. Make sure that the bass is natural and extended, that voices sound like real voices, and that the treble is not too dull or splashy. The more convincing the original sound, the more effective the use of TrueVerb in placing the sound in a natural acoustical setting.

Proper stereo width

Listen to whether the stereo width seems natural for the acoustic space and size of room you has placed the mix within. If you attempt to place a very widely panned left/right jazz recording into a 8000 cubic meter room, 20 meters away from the listener, then the sound will seem weird and unnaturally wide, with instruments 60 feet apart! The solution is either to choose a more appropriate distance and room size, or else to use a width control before TrueVerb processing to narrow the stereo image (such as the Waves S1-Stereo Imager). Then it will have a more natural stage width within the larger room you create in TrueVerb. In this example, don't be afraid that the reduction of stereo width will result in poor stereoism. Although the direct sound is narrowed, the simulated room created by the TrueVerb has a high degree of stereoism, and so will retain a sense of stereo wideness.

Natural sound+TrueVerb= maybe not

For instance, if the original recording already has some natural sense of distance and acoustical perspective because it is recorded with simple or purist microphone techniques in a live room or concert hall, then trying to add a TrueVerb effect may not be successful. The clash between the original acoustical perspective in the recording and the one created by the TrueVerb may not be what you are really looking for.

Using TrueVerb with previous Reverb effects

Because TrueVerb thru-processing places a recording within an acoustic space, its effect is subjectively quite different from most other reverberation processors on the market. Most of these produce a highly creative effect in its own right, but may be rather artificial. The design of other reverbs is distinctly different from that of the TrueVerb, but this leads to interesting possibilities when the two are combined.

The effect of adding TrueVerb thru-processing to a recording that already has other creative artificial reverb effects is to keep those other effects as they are: special musical and production effects. What happens is that they will be perceived as effects being produced electronically within the natural live room of the TrueVerb.

While no universal rule can be given for how the TrueVerb interacts with other reverbs, the different effects of TrueVerb and those other processes mean that there is plenty of experimental possibilities of combination. However, in general, the best way of doing this is to add the artificial creative reverb effects individually to individual musical parts in the mix, as is conventionally done in mixdown, and then to use TrueVerb thru-processing as a final stage after mixdown to place the whole mix into an acoustical space. If processing such as compression or stereo manipulation is applied to an overall mix, this is generally best done before the use of TrueVerb thru-processing, since this avoids changing the natural acoustical cues of the TrueVerb spatial simulation. Of course, you can be as 'illegal' as you wish!

In general, the use of TrueVerb processing to sweeten and place a final mix within an acoustical space will work best if the earlier processing with other reverb effects are of the highly creative kind, such as gated or nonlinear reverbs and various other kinds of exaggerated effects. Where trouble may sometimes be encountered is with reverb settings on other reverbs that attempt to simulate a natural acoustic. Although these attempts may be not very effective in themselves, they can introduce wrong acoustical cues to the ears that conflict with the cues given by the TrueVerb, and result in a less successful effect (if it's nature you are aiming for).

Chapter 7 - General TrueVerb Basics

When using TrueVerb for natural room simulation, ask yourself if the virtual room is appropriate to the music, and don't overdo or exaggerate the effect, unless needed! It helps to have experience of listening to live music in a variety of styles in good live performance conditions. When using the TrueVerb and selecting its settings, you then can focus on whether the sound is believable, and one that you can imagine having heard live.

The appropriate distance, room size and reverb (decay time) depend very much on both the music and the effect desired. It is impossible to give general rules here, but the following comments and guidelines may help.

How to begin

Typical starting-point settings of TrueVerb controls will place the absorption frequency around 5 kHz with an absorption of between -3 and -6 dB, and a HF damping factor of maybe 0.3x and frequency of maybe 7 kHz. The LF damping factor should be around 0.7x and frequency around 250 Hz. Depending on the room size and distance, the RevShelf may be set around -2.0dB but should be done by ear.

Unless you have good reason to do otherwise, use the default Balance of 0 dB, with the Link on, and Dimension at 3.0 (you are trying to recreate a 3-dimensional room!). After setting the distance, using the Balance control first and lightly is the best and most natural adjustment.

For the most natural effect, make sure that the Distance marker on the Time Response graph is to the left of the RoomSize marker. When you slide it to the right, past the RoomSize marker, the Distance marker turns red as a warning of a potentially unnatural effect: you've moved too far away, further than the room would allow in reality.

For music with an intimate quality, a small room size may be appropriate, but small room sizes will also have a higher degree of room coloration and maybe a less spacious effect. Also, the distances appropriate to small rooms may be quite small, only 3 or 4 meters.

RoomSize and Distance

Remember that there is no theoretical relationship between room size and reverberation times (Decay Time). A small but very live and reflective room (such as a tiled room used as an echo chamber) can have a long reverberation time, and a large but well-damped room can have a short reverberation time. However, most of the rooms you hear live music in tend to have a reverberation time that generally increases with room size. Small rooms often have reverberation times between 0.2 and 0.5 seconds; medium or club-size venues of a few hundred cubic meters may have a reverb time between 0.7 and 1.2 seconds; small concert venues, theaters and chamber recital rooms around 2000 to 4000 cubic meters may have reverberation times between maybe 1.1 and 1.8 seconds; and large concert halls may have reverberation times between 1.6 and 2.3 seconds. Churches, cathedrals and other grand architectural buildings may have very long reverberation times up to 5 or 6 seconds, and in very exceptional cases, reverberation times of up to 40 seconds have been encountered! Certainly, TrueVerb can be used for excellent ambient effects using a very long DecayTime.

Choosing the distance, room size and reverberation time appropriate to a music requires sympathy to the musical style and its needs. There is not necessarily a single correct answer either. For example, with a rock recording, you could choose to give a close on-stage sound in a small venue, setting up a distance of only 2 or 3 meters, and you may

choose a short reverberation time of 0.7 seconds to go with it. In this way, the reverb is not obvious but merely a part of the general acoustic picture. Or you might go for the sound of a good PA in very large venue of many thousands of cubic meters, with a distance from the PA speakers of 20 meters and quite a long reverberation time of say 1.6 seconds. This will give a very different altogether "larger" sound, but without the immediacy of the former choice. Both would be quite usable for different reasons.

Imagine the room

The important thing in making these choices is what works artistically, and also having a clear picture in your mind of the situation you are trying to create. A confusion in the goal is likely to result in a confused sound picture. There are obviously trade-offs that have to be made - it is unreasonable, for example, to expect simultaneously a very intimate sound with a lot of bright impact, and yet at the same time to expect the sound of a large PA a long way away in a large place. Both kinds of effect may have their attractions, but you have to make decisions about what aspects are appropriate to a given recording.

Large scale orchestral recordings will generally not sound best if placed too close or in a too-small room with too short a reverberation time. In general, orchestral reverberation times will be between 1.3 and 2.3 seconds, and distances of 10 to 20 meters may be appropriate, with room sizes from 4000 to 20000 cubic meters.

However, what is appropriate to a large romantic symphony orchestra will not be appropriate to a small chamber orchestra, such as used in authentic performances of music of the baroque or classical period. Here, smaller venues with slightly smaller distances and reverberation times may well prove more appropriate.

Chamber music is often heard in larger halls than it should be for the musical style. The reason is because of the need to finance concerts by having a large audience size. In using the TrueVerb with this kind of music, ask yourself if maybe a closer distance in a smaller room may not in fact be more appropriate musically for a more intimate effect. The same may apply to small jazz groups, which may often sound at their best in a intimate night-club size venue, listened to from quite a small distance of maybe 3 meters. Try and imagine the musicians actually playing in front of you without a PA at that distance. Is the effect convincing? Have a conversation with yourself and imagine the room and its sound.

Rock and some pop music provide some of the most interesting possibilities in the use of TrueVerb Thru processing, since rock music works both as an intimate "chamber" type music in small rehearsal rooms (when it can sound quite excellent with a clear intimate sound), and as a large-venue stadium music. Generally, rhythm and blues, early rock'n'roll and rockabilly is best placed in a small or medium size room, since these music were not conceived in terms of large venue performances. Most 'indie' rock is probably best also not placed in too large an acoustic space. This may be perfect for a virtual garage, with cheap rugs!

Whatever kind of music you are placing in a space with the TrueVerb, the best results will be obtained if you do not attempt to go for exaggerated effects, although these certainly can be achieved when required. Instead, processing that is relatively 'natural' often proves to be much more satisfying when listened to for longer periods. As always in processing recordings, what gives the biggest immediate impact may well not be what proves to be most satisfying or comfortable to listen to over the whole of a piece of music or an album. In that situation, a naturally unobtrusive effect may prove to be far better.

And often, the simple fact of the music sounding good and comfortable may have a bigger long-term impact on listeners than the initial "gee-whiz" impact of an initially spectacular sound. On the other hand, TrueVerb gives you a group of tools based on nature that you can use very unnaturally! And so it goes.

Chapter 8 - Setup Library

The setup library provides a range of settings for production, from classic rooms to ambient effects. Of course, you will create your own setups. The Load/Save function has been greatly improved in flexibility, and if you haven't read the Load/Save chapter, you should do so before reading any further.

About the setups

Documentation on each of the setups is not in this manual, but is in DocMaker™ and Word™ v5.1 format in the Setup Library 'Documentation' folder.

About the library

The setup library is divided into two parts for convenience. **Thru** and **Send** modes each have their own folders.

The primary difference has been explained at several points in this manual, but it is so important that we'll mention it again. Skip this paragraph if you just can't read it one more time. **Thru mode** means everything goes thru the TrueVerb, including the direct sound. In the world of pro audio, the **Send mode** has been more commonly used for many reverb processors, whose emulation of distance is a wet/dry control and maybe some early reflections. These are highly flexible and creative effects in their own right. It is simply that TrueVerb has brought a unique implementation to the market that requires the 'thru mode'.

Thru setups

Thru mode has templates that you can load to get started quickly, if you want to listen to the Setup Library setups; load the "Thru mode template" from the Templates folder in the Setup Library. Of course you'll have to import a stereo audio track, at the least.
Each file in the Thru and Send folders contains more than one setup. All the Halls are grouped together, all of the Rooms are together, and so forth, for easy comparison and access.

Send setups

Send mode also has templates for quick use. Load the "Send mode template" from the Templates folder in the Setup Library. Of course you'll have to import a stereo audio track, at the least.

The Send setups have a different configuration. In all of them, the Direct sound is switched out and the Link is off. The reason for this is that the Direct sound will come from the sending tracks. The Link is off so that you have fully manual control of everything. The **Balance** control has a much wider range, and the **PreDelay** will not be automatically changing with RoomSize. If you want the automatic PreDelay settings (which can be very convenient), just switch the **Link** on.

The **Link** button, when on.

PAZ

Psychoacoustic Analyzer

Table of Contents

Chapter 1 ..... Quick start 2

Chapter 2 ..... PAZ component plug-ins 3

Chapter 3 ..... The PAZ interfaces 4

Frequency window (RTA) 4

Stereo Position Display (SPD) 6

Level meters 7

Chapter 4 ..... For fast graphics 8

Chapter 1 - Quick start

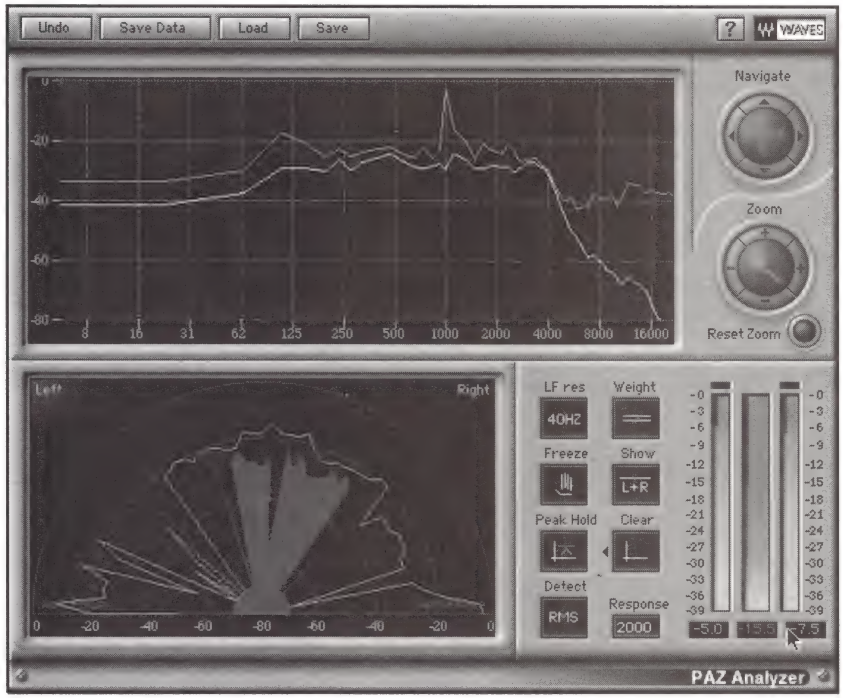

PAZ is so easy to use, you can almost work from the illustration shown above! However, for a more complete understanding of what the PAZ does and why we chose certain functions, please read this manual.

We suggest you use only the plug-in you need; this will give you the fastest graphic response time (if this is important to your use).

Chapter 2 - PAZ component plugs

On most host applications and platforms you can select just the meter you need. In the case of the PAZ, no DSP savings are accomplished, but graphic speed is increased by using smaller component plug-ins. The PAZ has several plug-in components in the menu. They are all mono and stereo compatible and do not change the audio at all, as they are fully bit-transparent (24bit fixed point or 32bit floating point clean).

Here are the current components as of this printing. New ones might be added in future updates, and will be noted in the plug-in Read Me files. Some examples for each component are listed:

- **PAZ-Analyzer** — This is the full plug-in with all 3 measurement functions (Frequency, Stereo Position, and Level); it has the biggest demand on the computer's graphic power.

- **PAZ-Position** — Shows only the Stereo Position Display with the controls related to it.

- **PAZ-Frequency**— This shows only the real-time frequency analysis with related controls.

- **PAZ-Meters**— Shows only the meters for Left, Right, and Sum. The Sum meter can be selected for Peak or RMS display.

Chapter 3 - The PAZ interfaces

Frequency window (RTA)

Here is the PAZ frequency analyzer graph. The screen shot is taken from the PAZ-Frequency component. In other components, the row of buttons below the graph may be arranged differently, but they are functionally identical.

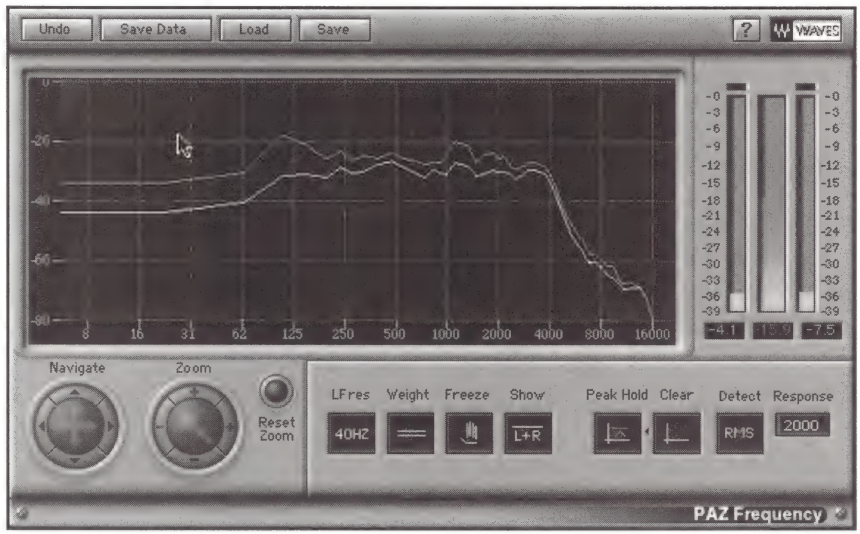

The PAZ frequency analysis is done in bands very similar to those of the ear. The default setting of the **LF resolution** control is 40Hz, which gives 52 bands, and most closely approximates the constant-Q critical frequency bands of the ear. The result is a good correlation between what we heard and what we see in the PAZ frequency analyzer.

For more detailed or technical uses, LF resolution below 250Hz can be set to 20 or 10Hz, providing up to 68 bands; the button's options are 40, 20, and 10Hz. Above 250Hz, the "engineering Q" (or width) of the bands are about 10.0, which as mentioned, are similar to the resolution of our hearing. The 40Hz LF default also closely matches the ears characteristics, plus has the best time response in the meter itself.

The **Peak/RMS** and **Response** controls let you choose the type of analysis and response time. In the **Peak** mode, peaks in each frequency band are displayed and the Response simply controls the release time. Under the RMS mode, the energy is averaged over time, and the Response value controls the length of this time.

Optimal time and frequency resolution in the PAZ is achieve by using wavelet techniques (as opposed to FFTs). This

lets each band update independently as fast as possible for its frequency resolution.

The real-time frequency analysis can be used for mono, dual channel, or total stereo energy by using the **LR** button, which switches between dual channel (**L/R**, independent channels) and total stereo energy (**L+R**). The L+R mode is not a simple summation of L and R signals, but of L and R energies. This more closely models our ears' responses to loudness of frequencies, as the gain and pan position of a sound are independent. The dual channel mode may be selected for precision analysis or troubleshooting.
The display can also be set for either of 3 weighting curves (dBA, dBB, or dBC), or unweighted. When unweighted, the absolute 'electrical energy' values are shown. Like in standard loudness meters, the weighting curves can be used for noise and loudness measurement applications that require them. Precise definitions of the monitoring levels for these weighting curves are beyond the scope of this manual, but in short, low frequencies offer little contribution to the overall loudness when monitoring at low sound levels. In other words, we are less sensitive to them. So an 'A' weighting is used which reduces the contribution of these LF to the meter reading. 'B' is for intermediate levels, and 'C' is nearly flat, for high level monitor loudness measurements. Unweighted is for precise, "what you see is what is there", measurements.

Option-dragging in the Frequency window will zoom the selected area. Click the **Reset Zoom** button to reset to maximum display of all frequencies and dynamic values.

The **crosshair cursor** can be shown by simply clicking once in the Frequency window. You can drag it around or click anywhere to show it. **Drag the crosshair to the upper corners** to hide the crosshair cursor.

At any time, the Frequency graph can be stopped by clicking the **Freeze** button. To release and return to realtime graphing, click the button again.

The maximum graphed value of the Frequency analysis can be shown by clicking the **Show/Hide** button. Maximum values are tracked even when in Hide mode. To reset the max graph line, click the **Clear** button.

A note on Pink Noise measurements: Pink noise is defined as noise that has constant energy inside any frequency band with a constant Q (i.e. frequency/bandwidth=a constant value). Pink noise is an important audio signal since it represents the average spectrum of a typical audio signal, and because the ears analyze the sound in approximately constant-Q bands. Because PAZ has almost constant-Q bands, a pink noise input will show up as a flat spectrum when no weighting is applied.

Stereo Position Display (SPD)

The SPD shows you the stereo energy in your mix, clearly and intuitively. This unique audio meter analyzes the distribution of loudness across the stereo stage, as well as detecting out-of-phase components. It differs from the common 'phase display' where left/right channels are displayed against each other on a oscilloscope screen.

For those familiar with the Waves S1 Stereo Imager, the Stereo Position Display will seem familiar, as both use a polar coordinates graph. Loudness is shown by length of the line and stereo position is shown by the angle (direction). The concentric lines are 20dB step indications.

In other words, a mono signal goes straight up (as will anything in the center of a stereo mix). A hard left or right sound will go 45 degrees to the respective sides. Most stereo mixes will have a fan-shaped display centered around the vertical line with occasional little spikes shooting out at greater angles.

Out-of-phase signals are shown between 45 and 90 degrees on each side. In the area beyond 60 degrees on each side have a red color in the background to indicate that signals in this range might be disturbing to the ears.

The controls

They are the same for the SPD as for the Frequency Display, except that when using the component Frequency plug-in, only the related controls are shown.

About Anti-phase

It is a bit tricky to define when a signal is shown by the Stereo Position Display or heard by the ears as "out-of-phase". This is not simply when the left and right channels happen to have opposite polarity at any given instant. Such a situation can happen if independent instruments are panned to different stereo positions, and it just happens, statistically, that the signal of each channel has opposite polarities. The meter will show a quick spike in the anti-phase area. This can also happen if a stereo delay or reverb is applied to an instrument and the delays arrive to each

channel at different times.
In these examples, the ears do not hear a disturbing 'out-of-phase' sound. For us to hear an out-of-phase sound, there generally needs to be a consistent opposite polarity over a period of time, which the meter will show.

The Stereo Position Display analysis takes this into account. Still there will be occasional fast spikes in the anti-phase region, and you can ignore them if they are not consistent, especially if you know you are using stereo delay or similar effects. You might be really bothered if the energy in the out-of-phase region is consistently significant compared with the energy of the 'in-phase' region.

It is advised that you take the time to explore how the SPD looks when you intentionally flip the phase on one channel of a stereo signal. (You can do this by using the Q1 or one of the S1 components in front of the SPD).

Level meters

Peak and **Loudness** meters are shown beside each other; when in stereo, the L and R channels have separate Peak meters, with a summed Loudness meter which has 6dB subtracted from it (L+R-6dB). Although the true energy the ears are exposed to by stereo signals is indeed the sum of L and R energies, by subtracting 6dB from the loudness meter, it's easier to evaluate the Peak/RMS relations of the signal. This is because summed left/right signlas can have peaks as much as 6dB higher than either channel separately.

You can select **Peak** or **RMS** measurement for the Loudness meter, which is also affected by the selected **weighting curve** (unweighted, dBA, dBB, dBC).

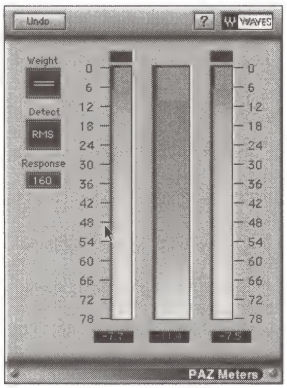

For those of you who own the L1-UltraMaximizer or the L2 hardware Ultramaximizer, it will be striking to see what it does to the Peak-to-RMS measurements of your audio.

Each of the 3 meters have **peak-hold values** at the bottom of the meter. Each can be **independently reset** by simply clicking directly on the peak hold value.

Clip lights are only provided on the full plug-in, not in the component Meter. They can be reset by clicking directly on the light.

Chapter 4 - For faster graphics

Fast graphic hints

Lower the number of colors in your Monitors control panel (256 instead of thousands or millions).

Use only the component of the Analyzer that you really need. For instance, if all you need is the frequency analysis, then use only the "PAZ-frequency" plug-in. The full plug-in with all 3 meters will have the heaviest demand on the host computer.

Under ProTools, if you click and hold on any control, such as the Release control, the Analyzer graphic will take all the computer "idle time" for graphic updates. This will freeze all Pro Tools meters, but will not affect any functions of Pro Tools (such as playback or recording).

De-Esser

Table of Contents

Chapter 1 ..... Introduction 2

Chapter 2 ..... Quick Start 3

Chapter 3 ..... Product Overview 4

Chapter 4 ..... DeEsser Controls 5

Chapter 5 ..... DeEsser Internal Design 7

Chapter 6 ..... FAQs 8

Chapter 7 ..... Factory Preset Information 9

Chapter 1 — Introduction

Thank you for choosing the Waves DeEsser – a versatile audio plug-in for selective and creative compression of high frequency 'ess' sounds in recordings.

The following pages explain how to use DeEsser.

Chapter 2 - Quick Start

1. Let us assume we want to reduce the 'esses' in a male vocal track. Click on the Load button and select the "Male Ess" factory preset. This sets the SideChain frequency to 4500 Hz, the frequency around which typical male voice 'esses' lie.

2. Drag the threshold down until the 'esses' are sufficiently attenuated. Attenuation takes place when the threshold slider is below the peak of the blue-bar meter (energy detector). The red-bar meter shows the instantaneous gain reduction applied to the audio (in dB). An infinite-peak-hold feature displays the highest level of attenuation (click to reset).

3. To fine-tune DeEsser, click the Monitor button to SideChain mode. Now you can listen to the SideChain. Choose which SideChain filter mode you require from either HighPass or BandPass. The HighPass filter mode is suitable for attenuating a full range of 'ess' sounds. The BandPass filter mode is suitable for isolating and attenuating a specific narrow band of high frequencies.

4. Simply listen to the sound in the SideChain whilst adjusting Freq until you hear just the 'esses' for reduction and little else of the remaining signal (that is, not too much of the voice itself).

Chapter 3 — Overview

We are all familiar with the harsh 'ess' sounds orators or singers can produce during a powerful delivery close to a microphone. Almost any vocal recording will contain 'ess' sounds, whether from a strong vocal delivery, bad recording, speech impediments or simply many 'ess' words spoken together. Wind instruments and other musical instruments can also create shrill high-pitched noises.

Sound engineers need to control these harsh 'ess' sounds in most recordings. Here is a graph of a strong ess sound from a vocal track (graph from Waves PAZ). You can see the energy strongly centered about 12kHz (rather high for an ess sound). DeEsser is an audio plug-in for attenuating higher pitched frequencies such as those found in 'ess', 'shh' and 'chh' sounds.

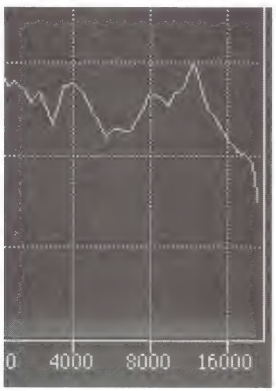

Featuring very sharp filters in the SideChain and the choice of either Wideband or Split audio paths, DeEsser is suitable for processing full mixes, solo vocal-takes and instrumentals. Gentle DeEssing on complete mixes produces great sounding results.

DeEsser enables sound engineers to creatively compress 'ess' sounds. Here is the same section shown previously, after de-essing from 3kHz and higher:

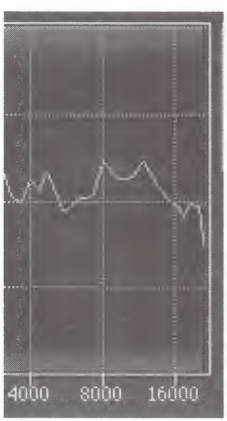

DeEsser graphic interface (Mac shown; PC is same)

DeEsser Controls

Kindly refer to the WaveSystem Manual for a full explanation of our standard controls.

Thresh

Adjusts the level above which attenuation takes place in the SideChain. To adjust the threshold, click and drag the slider down until the 'esses' are sufficiently reduced. NOTE: lowering the threshold too much will result in a muffled sound (i.e. too much attenuation).

Audio

Toggles between Wideband and Split compression modes. In Wideband mode, attenuation is applied to the entire audio band. In Split mode (i.e. the audio is split into low and high frequency signals), attenuation is applied only to the HighPass signals.

Freq

Sets the frequency level in the SideChain filter using a value window control (see WaveSystem Manual for further details).

SideChain

Toggles between a HighPass and BandPass filter. The HighPass filter 'looks at' all the frequencies above the set frequency, the BandPass at a narrow band around the set frequency.

Monitor

Use this to listen to just signals in the SideChain. When in SideChain mode, a yellow light below the Monitor button switches on.

DeEsser Displays

Energy Detector

Energy of the SideChain (dB below 0 dBFS). The energy bar-meter and threshold-slider are aligned for easy adjustments and visual feedback.

Attn

Instantaneous gain reduction (attenuation) applied to the audio signal (in dBs).

Output

An infinite peak hold feature shows the highest output level. Click to reset.

Chapter 5 - DeEsser internal design

DeEsser internal design

Signals entering DeEsser are fed into the audio path and SideChain.

In the SideChain, audio is filtered and translated into attenuation - a gain reduction of zero or more dB. This attenuation is then applied to the audio path, giving the output.

SideChain

The SideChain consists of a filter, an energy detector and a compressor. The filter has two modes of operation: HighPass and BandPass. HighPass mode is useful for attenuating several different 'ess' sounds. BandPass mode is more suitable for attenuating a specific high frequency. The chosen filter's frequency is set by the Freq. control. NOTE: This filter is only related to the SideChain, not the audio, which can be either Wideband or Split.

The energy in the filtered signal is measured and translated into attenuation by a hard knee compressor according to the threshold set by the user. When the SideChain energy is below threshold, the compressor generates no attenuation. When energy exceeds the threshold, attenuation is generated – more energy above the threshold means more attenuation.

Audio path

In Split mode, audio is split into high and low frequencies. Attenuation generated by the SideChain is only applied to the high frequencies, while low frequencies remain unaffected. In Wideband mode, attenuation is applied to the entire audio.

Chapter 6 — Frequently Asked Questions

When should the SideChain be in BandPass mode?

In general, the BandPass filter is used to isolate one type of 'ess'. The BandPass filter 'looks' at a narrow range of frequencies. The 'ess' sound in "Sweat" carries most of it's energy around the 6kHz mark (male voice). In "Sugar" most energy lies around the 4kHz mark (also male voice). Use the BandPass filter if you only want to reduce one 'ess' and not the other.

When should the SideChain be in HighPass mode?

If you have several different types of 'ess' noises you wish to reduce, the HighPass filter is recommended. DeEsser responds better in HighPass mode when applied to, for example, a vocal track containing a variety of 'ess' frequencies and vocal inconsistencies. NOTE: the HighPass filter increases the chance of compressing other high frequency sounds such as strings and high pitched instruments.

Why can the audio be either Wideband or Split, and how do I choose?

Most compressors containing a sidechain component can be used as de-essers. However, they require an external equalizer to select the frequency range for reducing. Such compressors have Wideband audio passing thru them and thus compress all frequencies equally.

This is the most traditional approach for purist or vocal-only applications. For powerful 'esses' and the more modern vocal styles, splitting the audio into 2 pieces is more effective and results in fewer side effects on frequencies outside the 'ess' range. If you must de-ess a mixed track of instruments and vocals, the Split method is nearly essential.

Choose the Split mode for the most accurate and strong de-essing and for mixed tracks.

Choose Wideband mode for gentle voice-track-only de-essing and when you do not want to have a crossover in the processing path.

Why would I want to monitor the SideChain?

Only to make it easier to adjust the Freq. You can "dial in" the best Freq value by simply monitoring the SideChain. Simply listen to the sound as you adjust the Freq until you believe that you are hearing what you want to de-ess and as little as possible of the remaining signal (that is, not too much of the voice itself).

Is it possible to over DeEss?

Absolutely! If your singer starts sounding as if you have removed their two front teeth, then the Threshold is too low! Adjusting the amount of de-essing is subjective - just listen and use your judgement. When the voice sounds more natural and the 'esses' are not overpowering, that is enough!

Chapter 7 — Factory Preset guide

The factory presets appear under the Load button menu (as described in the WaveSystem manual).

General settings are provided for male and female voices, although they can be used for instruments as well. Setups are provided as convenient starting points. The difference between each setup is the starting frequency. All presets set DeEsser to Split mode and use the BandPass SideChain filter mode.

Kindly note that "shh" settings are also relevant for 'ch', 'th', and hard consonants, such as 't', 'd', and 'k'.

Default

When you first run DeEsser, the default settings are suitable for almost any speech or song recording.

The default frequency setting is 5506Hz, which will attenuate down to about 3kHz, the main attenuation taking place from 5kHz up.

The default SideChain filter mode is HighPass for general de-essing.

The default audio setting is Split mode, meaning the audio path is split into two parts - only high-frequency ranges receiving attenuation. This mode is more commonly used, whether for voice or for full-mixes.

As with all presets, simply adjust the Threshold until some attenuation takes place, then adjust the Freq to affect the desired range.

Male Ess

The Freq is set to 4500, as most males have lower 'ess' frequencies than females. An exception may be older men who sometimes have rather strong and rather high-frequency whistling.

Therefore this setup covers the lower range of possible 'esses' for the average male.

Male Shh

Most 'shh' sounds are nearly one octave lower than 'ess' sounds. However, if the Freq is set too low, then too much of the sound is affected. Therefore, a Freq setting of 3385Hz is set for the Male Shh. DeEsser will 'reach' down to about 1700Hz, but with lesser and lesser maximum gain reduction.

Female Ess

Female speech and singing generally has higher, sometimes much higher-pitched 'ess' sounds than male. This setup has a Freq of 6779.

As with all the factory setups, the frequency settings were determined by averaging a large number of voices. The Freq may have to be adjusted a large amount to accommodate a certain voice - with some female voices, the 'ess' sounds can range as high as 9kHz center points!

Female Shh

The frequency setting of 5077 generally covers the female 'shh' range which, although lower than the 'esses', is not as harsh or as low as the male 'shh'.

Generally, female 'esses' and 'shh' sounds vary more in frequency than males. In other words the various frequencies of the 'esses' and 'shh' sounds of a female may have a wider frequency range than those of a male. Thus you may find that using the SideChain filter in HighPass mode more responsive.

Conversely, in BandPass mode only the specified frequency band is measured.

To summarize, if a singer has 'esses' spanning a large range of frequencies, try the SideChain filter-mode set to HighPass.

SuperTap

Table of Contents

Chapter 1 ..... Introduction 2

Chapter 2 ..... Quick Start 3

Chapter 3 ..... What is a Multi-tap Delay? 5

What is a delay line? 5

Modulation 5

Chapter 4 ..... Tips and Hint 6

Read about WaveSystem 6

Work with the presets 6

Chapter 5 ..... Controls & Displays 7

Pan graph 7

Tempo 7

Modulator 8

Output 9

Direct 9

Delay line (taps) 9

EQ section 11

Feedback 12

Norm mode 13

Tap Feedback mode 13

Chapter 6 ..... Factory Presets 15

Chapter 1 - Introduction

Thanks for buying Waves processors.

SuperTap is a six-tap delay plug-in that can be used to produce a broad range of multi-tap delay effects, including analog and tape delay emulation, chorusing, rhythmic looping, and more. It also includes modulator and feedback sections.

The following pages explain how to use SuperTap. Be aware that this processor, just like its analog equivalents, is easily capable of overmodulation and positive gain feedback loops (i.e., increasingly louder sounds going well into distortion). There's no way to give you complete freedom and protect you from getting such sounds, so Be Prepared.

Chapter 2 - Quick start

At the most basic, there are two types of SuperTap processing: one with modulation, or one with longer taps. You can also choose between 2 or 6 taps for either type, as well as mono or stereo. Please note that the TDM and Native Power Pack products come with a 2-tap-only version of SuperTap; the full version is part of either the Pro-FX Plus or Gold products.

In your menu, you'll see four different component plugs, or "sub-plugins" of SuperTap: 6-taps (with and without Modulation) and 2-taps (with and without Modulation). The 2-tap takes less power and memory. The Modulation versions have shorter delay times (1.2 seconds); the components without modulation have the longest delay times (6 seconds).

Note: The components without modulation are not available in native applications. Maximum delay time may vary from application to application

Launch one of the SuperTap components from your plug-in menu. It will open with the default settings. The default setting of any of the components has all the taps turned on, without Feedback. The factory presets (located in the Load menu) are a great place to start, but take a little time to read at least this page! There's lots to know about this processor that will save you time.

1. Each tap's controls are arranged in a row from left to right. The on/off button is on the left, going across to the EQ section. Notice that when you turn a tap on, it illuminates in all three buttons: the on button itself, the delay slider, and the Pan Graph (for stereo plugins).

2. Drag the tap marker to set the delay time for each tap.

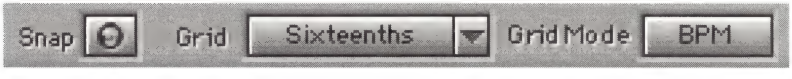

3. Drag a Pan Graph marker to change Gain and stereo position (rotation) at the same time, or adjust the Gain and Rotation controls separately for each tap. (Panning is only for mono signals, Rotation works with mono and stereo, so that is why it is here).

4. On the right side is the EQ/Filter section, with on/off control, 5 filter types, plus gain control for the bell and shelving filters. You can filter each tap separately for a very large range of creative rhythm effects from a single sound source.

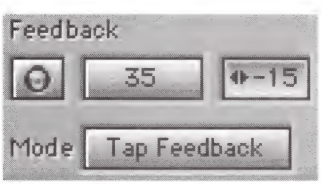

5. The feedback section in the lower left corner is worthy of a warning! Using feedback can produce increasingly louder sounds, well into distortion, just as the analog equivalents can do. In order to give you great flexibility, there is also the small danger of very loud feedback sounds, so it is best to use low feedback values and adjust it slowly.

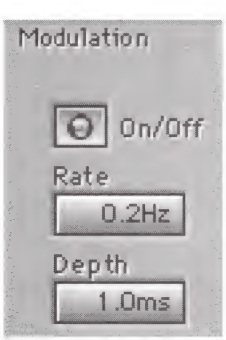

6. Modulating the taps produces variation in the delay times, which changes the pitch from each tap. You can control the depth (in milliseconds) and the rate of variation.

Chapter 3 - What is a Multi-tap Delay?

Audio delays have widespread use in studios and sound design and have been used for quite some time, dating from the earliest recording devices. There have been many versions, including analog "bucket brigade" types, digital, acoustic, tape, and more, each with their own flavor and character. The SuperTap doesn't try to directly emulate all of them, but you have a huge set of tools to work with to do exactly that, or create completely new effects only available in digital software tools.

What is a delay line?

A delay line simply plays a copy of the original signal following a delay of a certain amount of time. When you add this delayed copy to the original, some extraordinary effects can be created. When you play several such delayed signals along with the original where each signal is delayed by a different amount, you have a multi-tap delay.

Therefore, a multi-tap delay is a single delay with multiple outputs (taps) along the delay. In SuperTap, the delay is a maximum of 6 seconds (without modulation), and the taps can be placed anywhere along that delay. These taps are "feed forward", which means they are all mixed together at the output. A "feedback" control is also provided which routes audio back to the input for looping and complex rhythmic effects.

Digital reverbs operate using a similar principle. By adding mirrored copies of the original signal to itself, and by reducing the gain each time, the effect of a disappearing echo in a large space can be reproduced. However, reverbs can have millions of delayed repetitions with great complexity in relationships (volume, phase, frequency response, direction).

Modulation

Audio modulation is the variation of one or more aspects of a sound. There are many kinds of modulation (frequency, amplitude, phase), but in SuperTap we use only Frequency modulation. A low-frequency oscillator (20Hz and below) very slightly changes the delay times of each of the taps, producing changes in the pitch of each output. This is ideal for chorusing and other effects where varying pitch is desired.

Chapter 4 - Tips and hints

Read about WaveSystem

In one sense, SuperTap is very simple, but can be so complex that it is best to work with it to learn what is going on. However, there are some tips that can help you work more effectively. One is to take the time to read the WaveSystem manual; this is mentioned a few times in this manual, and for very good reason. Waves is the world leader in software-based processors, and we've developed highly intuitive and powerful interfaces. But some features are unique to Waves, so unless you read about them, you may never find them.

One example: you'll notice there is no wet/dry control on SuperTap. It is easier to control the Gain of each tap separately. But what if you want to control all of them simultaneously?

It's quite easy. Simply click and drag across all Gain controls to select them with a selection rectangle, then

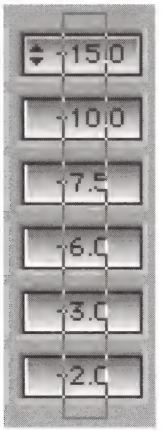

click and drag on any one of them to change all of the Gain controls at the same time, as shown here: You can use the arrow keys to tweak values, or enter values directly from the keypad, and much more. So, read the WaveSystem manual for powerful control that applies to all Waves plug-ins.

Work with the presets

Each of them were developed to illustrate a specific aspect of SuperTap, and although hardly exhaustive of the processor's capabilities, they are great places to start and tweak. By using them and simply adjusting the Tempo to match your production, you can save a lot of time. We encourage you to post any settings you create to our web site.

Chapter 5 - Controls & Displays

Pan Graph

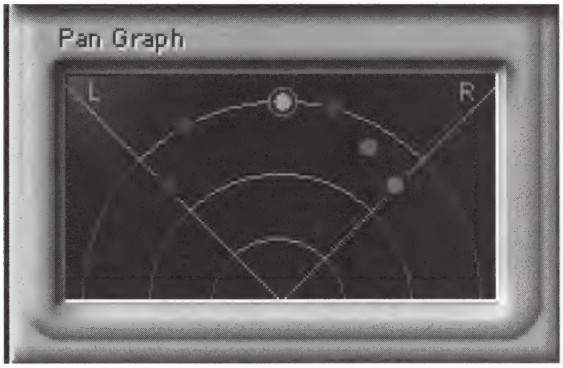

The Pan Graph shows the stereo location of each tap in a stereo field. (Since the SuperTap also works with stereo inputs, the correct term is Rotation, which applies equally to mono or stereo inputs). You can click and drag each tap's marker inside this display to adjust its pan (Rotation) position and mix (Gain) values. Drag a marker from side to side to change its Rotation value (from -45 to +45 degrees), and up and down to raise or lower its output gain (from 0.0 to -24dB).

Tempo

You can use the triangular push buttons arrows to raise or lower the tempo value, or directly enter a number, or drag on the BPM value display with the mouse. One very important thing to remember about the Tempo control: Beats Per Minute always refers to the quarter note. No matter what Grid mode you are in (i.e., whether sixteenths or eighths, etc.) the BPM refers to the quarter note. The tempo is displayed in two forms: BPM (beats per minute) and ms (milliseconds).

The Mode button changes the function of the Tap Pad between Tempo and Pattern, depending on the desired mode.

Tempo Mode

If you don't know the exact tempo (speed) of the track you want to edit, but can listen to it, use tempo mode. It allows you to set the tempo rate by 'tapping' the pad with mouse clicks. The tempo is calculated automatically. In fact, SuperTap continually calculates the average tempo while displaying it in the BPM and ms displays. To get the most accurate average reading, tapping for a longer period is preferable. For example, if you tap reasonably consistently for awhile, (even up to two minutes!), the setting should reach a fixed value without changes despite the erroneous tap here and there.

Pattern Mode

If you want to enter your own sequence of delay times for each tap, use Pattern mode. Click your mouse on the Tap pad using your chosen pattern. You'll notice each tap takes the same delay as the sequence you tap. Each tap 'jumps' to position in the tap grid according to the sequence. This takes a bit of practice but can be used to quickly set a complex rhythm while, for example, listening to the song you're working with.

Modulator

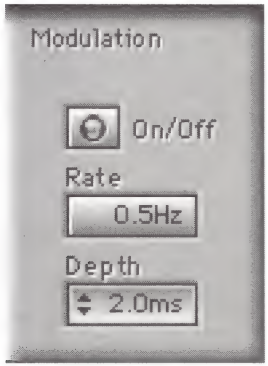

To switch on the Modulator, click on the On/Off switch. The modulator only affects the active taps, not the direct.

You can use the Rate button to adjust the oscillation rate. Its range is from 0.0 to 20Hz. The oscillator sweeps the delay of each tap according to the Depth setting. In other words, if the Depth setting is 250ms (its range is from 0 to 500ms) and a particular tap's delay is 700ms, the delay of that tap will sweep from a maximum of 950ms to a minimum of 450ms.

Output Gain

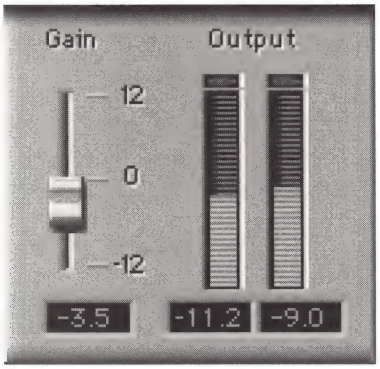

Next to the Modulator are the output gain controls and peak meters. Click and drag the fader to raise or lower the output gain (range is from -12.0 to +12.0dB); the value is displayed below the fader. There are clip lights located above the level meters (click to reset). Below the meters is a numerical display showing the exact value in dB of the highest peak level since the last reset, for both left and right channels.

Direct

The On/Off button turns the direct tap ("dry") on or off. When you launch SuperTap, the Direct tap is on by default.

You can change Gain and Rotation of the dry signal just as you do for each of the taps, and as well, you can click and drag the Direct marker in the Pan Graph.

Delay lines (taps)

Each tap has its own set of controls arranged in a horizontal row, as seen here.

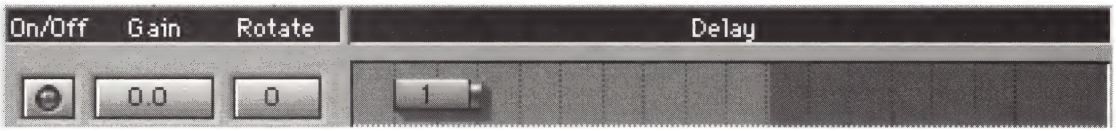

On/Off

This button, resembling an LED, can be used to switch a tap on or off. When a tap is switched on, a marker with a matching color appears in the Pan graph and a light on the side of a tap's marker (in the Grid display) illuminates. You can adjust the Gain and Rotate parameters of each tap using the controls described below:

Gain

Gain is adjustable from 0 to -24dB FS, in 0.1dB steps.

Rotate

Rotation is adjustable from -45 to +45 (in degrees). Rotation is a type of panning, but works for stereo as well as mono. For further information, the Waves S1 Stereo Imager manual has much more information about Rotation. In fact, it is the original "pan" control invented by the man who filed the original patent for stereo, Alan Blumlein. Suffice it to say that it is much better than a pan control.

Tap Delay Grid

Each tap has its own slider located in the display Grid. Each slider shows a value. When in Beats Per Minute, the slider shows how many beats that delay is set to. When in Milliseconds (ms), it shows the value in milliseconds.

Keep in mind that if you have Sixteenths selected for the grid, the number shown in each tap is how many sixteenths that tap is, so if you have "4", that is 4 sixteenth notes, which is equal to a quarter note. Beats per minute *always* refers to the quarter note, no matter the Grid setting.

If you are working with Snap turned off, then the numbers also show whether they are slower or faster than the closest beat by the plus (+) or minus (-) sign in front of the number. If the slider is perfectly aligned with the grid, then the number has no sign.

When a tap is active, a color-coordinated light on the right of the slider illuminates. A slider displays the delay value for that tap in either milliseconds or BPM, depending which mode the grid is set to. Click on the Grid Mode button to change between BPM or ms. To adjust the delay value for a tap, you can click and drag its slider from side to side inside the Grid display, or use the arrow keys.

In addition, a Snap-to-grid button lets you control the "quantization" of the taps.

The grid can be displayed using the divisions selected in the popup menu below the grid:

- Free
- Quarters
- Quarter Triplets
- Eighths
- Eighth triplets
- Sixteenths

In order to simplify the grid and numbering schemes, the grid was limited to these selections, but the BPM has a very wide range of adjustment (40 to 1200bpm), so if you need to work on "half-notes" for instance, you can simply cut the BPM in half and use the quarter note settings, etc. For complex rhythms such as 5's, 11's, etc., we leave that to you and your calculator to work in the "Free" mode!

EQ section

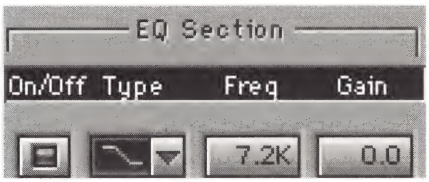

The filters on each tap are the same components from Waves highly acclaimed Q10.

Each tap (and the Feedback line) has its own filter. If you want to hear to the results of changing a filter's parameters on one tap, turn off any other taps and the direct tap.

Each filter's range is from 0.1kHz (100 Hz) to 20kHz and there are six different filter curves:

1. Bell 1 (wider)
2. Bell 2 (narrower)
3. Hi Shelf
4. Low Shelf
5. High Pass ("Low cut")
6. Low Pass ("High cut")

Tap Filter Controls

On/Off

Click on this LED-button to switch the filter on or off. When it's on, the button illuminates. See how easy this is?

Type

This button lets you choose a filter curve. Click on it to toggle through the six different filter curves listed above. Alternatively, click on the small arrow to select from a popup menu.

Freq

This control button allows you to adjust the filter's center frequency. Click and drag up and down on the Freq button, or click it and type in a numerical value. Each filter has a range from 0.1kHz to 20kHz.

Gain

Use the Gain button to boost or lower the gain of the Shelf or Bell filters (it does not affect the Pass filters, even though you can change the value).

It is possible to create a little rhythm section from one percussive strike using a different filter on each tap. For example, you can take a single bass kick, and produce 6 different tones or timbres with the 6tap processors.

Simply set each tap to have very different EQ settings, such as High pass on Tap 1 at 1.0kHz, Bell2 onTap 2 at 7.5kHz and Gain of 8dB, and low shelf on Tap 3, 200Hz, and with Gain set to -12 (for a thin sound). Then simply set the pattern with the Tap pad or manually.

Combined with Feedback and modulation, the possibilities for variations in looping are limited only by the amount of time you have to sit and work.

Feedback

As you were warned earlier, be careful with the Feedback! It is very easy to create very loud sounds that continually get louder and louder, and distortion is completely allowed in this plug-in. Otherwise we would have to make this processor so "safe" that many creative areas would have been eliminated.

The controls are:

- Mode
- Feedback on/off
- Gain
- Rotation
- Feedback EQ/filter section, except with different filters than the taps.

Mode

This is the most critical button in the Feedback section to understand.

In Norm(al) mode, the feedback simply sends each tap back into the input again. You'll notice that the Feedback delay slider is grayed out and can't be moved while in the Normal mode.

NOTE: while in the Normal mode you have the greatest risk of having positive gain feedback, that is, a sound that increases loudness rapidly into distortion. You must be careful about this, but at the same time, understand there are many creative areas that this allows. If multiple taps align onto the same time and you also have high Feedback settings, then positive gain is rather likely. In other words, it will get very loud very fast, so keep a hand on your monitor level until you "learn" what will happen.

In Tap Feedback mode, the feedback has its own delay time (like a "7th tap"), so if you set a rhythm using a few taps and set the Feedback delay time to the length of a measure, the entire rhythm will repeat. A gain setting of 50% means -6dB (half the energy), so the loop will fade out, where a setting of 100% will cycle a single input endlessly. It is in this Tap Feedback mode that the EQ/Filter settings give great power in making subtractive synthesis looping (filtering each tap differently).

On the other hand, if you set the feedback delay to a short length, and all other taps longer than it, then you'll hear each tap, followed by the feedback loop with each tap's sound overlaying into the short loop. This creates very complex timing issues that are better heard than described.

Feedback on/off

As Bob Clearmountain is reputed to have once said, if you have trouble understanding this type of button, maybe you should consider a different line of work.

Gain

Controls how much signal is fed back into the input. Warning: values above 50 are always delicate, so start with low Feedback Gain values and change it slowly.

Rotation

Being able to "pan" a stereo signal is at its coolest when used in the Feedback loop. For a mono signal, Rotation is very much like a pan control without a pan law (no volume dip in the center), but for stereo inputs, Rotation provides amazing flexibility and unexpected effects. They are unexpected largely because the world has little experience with Rotation, which moves an entire stereo image without collapsing the image width (for more info read the S1-Stereo Imager manual, available at our web site if you don't own the S1).

By taking a signal and rotating it all the way to -45 or 45 degrees (hard left or hard right), the entire stereo signal will be shifted dramatically each time it comes thru the feedback loop. This creates antiphase signals and very interesting stereo placement on looped signals.

For an example of Rotation on Feedback:

- Set just one tap for a 500ms delay.
- Set Feedback in the Tap Feedback mode with 50% Gain.
- Set Feedback delay to 1000ms.
- Set Rotation to -45.
- Play a small section of sound into SuperTap and listen as the sound rotates thru the stereo field (and even "outside" of it).

EQ/filter section

Only 3 filter types are in the Feedback EQ section:

- Low Pass
- Low Shelf
- Low Decay

The low shelf gives enough control to act as a highpass, and leaving out high shelves and high pass filters eliminates the possibility of strong boosts in treble. This keeps positive feedback overloads from producing huge amounts of high frequency energy. In other words, it was "safer" for us to make it this way.

Low pass and low shelf are self-explanatory, but the Low Decay is the filter type that is different from any other tap. It is simply a very gentle high shelf that has only cut (negative gain) values, so that you can make each repetition be a little duller (less treble), very much like tape or analog delays, or vintage delays that had low sampling rates. Good settings to start with are about -3dB.

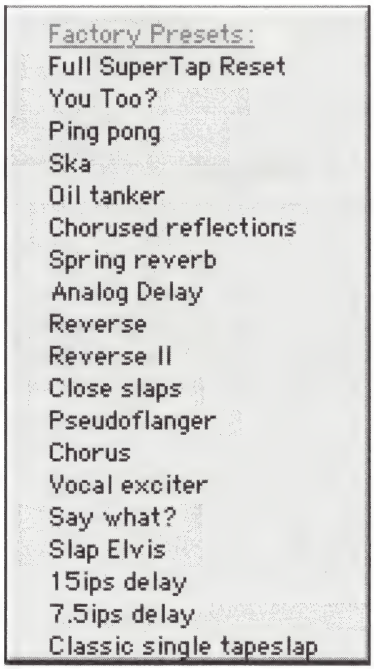

Each preset provides a good starting place for some fun. Here are brief descriptions of the presets.

Full SuperTap Reset

This is the default setting for all SuperTap plug-in components (mono/stereo, 2/6 taps, with or without modulation). It is set for fairly short delays of all available taps, with increasing stereo spread on each successive tap, as well as switching sides between L and R on each tap.

You Too?

A simple 2-tap cloner than duplicates the input with some Feedback for an audio house of mirrors. Quite good for a lead instrument, although you'll probably need to lower the gain of the taps some.

Ping Pong

Simple back-and-forth, with some HF dulling for better effect.

Ska

You know, the mighty mighty beat. The big big vocal. One Step Beyond. The Specials, in all their glory, or at least as we remember it.

Oil Tanker

Bad reflections, like a cheap house of too many mirrors, with closely spaced shaky reflections (sounds even better with the modulation versions).

Chorused reflections

A pleasant scattered reflection series with subtle pitch shifts. By putting the Tempo mode into Pattern mode and altering the Tempo, you can control the density (short or long) of the reflections. Very nice when blended with a deep digital reverb.

Spring reverb

Not as bad as the real thing, but with just as many distracting qualities. At least it's a start to getting toward the spring sound, and has the added feature and advantage of not making a huge noise if you bang on the top of the box.

Analog delay

Notice the filtering that heavily removes the high frequencies from the delays. By playing with these filters, even heavier dull emulations can become dark spiraling blurs. Very nice for cloudy, moody rock (which always seems to be popular with youth, no matter the age).

Reverse, Reverse II

By carefully setting the Gain of each tap, a pseudo-reverse effect is established, not much unlike a reverse-tape repeating echo (by flipping the reels over). By simply offsetting the track to the value of the last tap, you can have true pre-echo.

Close slaps

Nice for lead instruments or singers, as well as musicians. Old joke. The close spacing of the slaps can produce cancellation in the low frequencies, so if you wish to use this setting on a bass instrument, you might want to use a high pass filter on all the taps so that those frequencies don't suffer from unequal frequency response. If you don't understand this, just try it anyway, by using High Pass filters at about 500Hz or so on all the taps. The proof is in the listening.

Pseudoflanger

Why not use the MetaFlanger? It will do a better job, but if you need to switch to a setup (and since no DAWs allow you to automate the plugins being inserted or removed yet), this gives you a nice fake start on a flanging type of sound.

Chorus

Lots more to be explored here; we've only started you at the edge of a vast range of possible settings for chorusing and double/triple/n-pling the sound. Since the Rate and Depth don't have linking, you'll have to

carefully adjust both settings to get just the right amount of pitch variation. Similarly to the Pseudoflanger, closely spaced delays will cause cancellation of frequencies, so if you are using this on a source with a wide range of frequencies (especially bass or low notes), then high pass filters on each tap will help to control the unevenness of the sound. Better to have a solid bass in the Direct than to have unequal bass everywhere.

Vocal exciter

Another area with lots more to be explored. This is just a start, using 2 taps and some strong EQ to add an edge. By no means is it a true exciter, but it certainly puts an edge on the input.

Say what?

Thissiskindofhardtounderstandstandbecauseeverythingthing is runninginging togetherehter. A deliberately confusing set of time delays that makes comprehension rather challenging.

Slap Elvis

He'd probably have liked it, but this is really not about any physical issues, just giving the rockabilly vocal slap that seems to emulate that period of history, and in stereo too. Swing dance needs this type of effect if you blur the line between Brian Setzer and the King.

15ips delay

True emulation of the Ampex 440, whose distance between the record and play heads is 2.508 inches (thanks to John French of JRF Magnetic Sciences), which gives 172ms delay. If you fed the output back to the input, then you get the repeating taps we have set in this setup. Changing the EQ settings can provide something you never had with the old-style "turn up the tape output on the board" method: independent EQ on each repeat! This is as nice as having a holodeck Ampex 440 with a cloning subroutine.

7.5ips

Everything we said about the 15ips delay, but at half the tape speed, i.e., 344ms delay times. Very very much a reminiscent sound (to those who were there and can remember) of the types of delays during the period when the Ampex 440 was found in nearly every studio.

Classic single tapeslap

Simply a 1-tap version of the 15ips delay setup. This is probably more realistic a delay setting than the Slap Elvis, and is in mono.

IDR
(Increased Digital Resolution)

Table of Contents

Chapter 1 ..... Looking for a Quick start? 2

Chapter 2 ..... Waves IDR basics 3

Chapter 3 ..... About Dither & Noise Shaping 4

Chapter 4 ..... Recommended IDR settings 8

Chapter 1 - Looking for a Quick start?

There is a small section on using IDR inside the WaveSystem manual. Please refer to it for complete detail on when to use IDR (within the L1, for example) and how to use the small IDR plug-in.

The following section is intended as a technical overview of Waves IDR, as well as general information about wordlength reduction concepts, dithering, noise-shaping, quantization noise, and more.

Chapter 2 - Waves IDR basics

IDR is Waves' proprietary noise-shaped dithering system co-developed by Waves with Michael Gerzon, a world authority in psychoacoustics. The IDR component in Waves processors can be used to maximize the digital resolution of the resulting audio. IDR's design is a result of Mr. Gerzon's long-term research dating back to 1982 with many of the other leading experts in digital resolution enhancement technologies.

In the industry, IDR and such systems are interchangeably referred to as wordlength reduction systems, bit-depth reduction systems, dithering systems, and "proprietary processes". They are sometimes incorrectly referred to as downsampling, which relates only to sample rate. The end result (or goal) is the same, no matter the name: to capture more detail from a long wordlength (24bits or more) into a shorter wordlength output (such as 16bits).

IDR controls

The Waves IDR system is implemented using the following three elements:

1. Quantize (output wordlength)
2. Dither type
3. Noise shaping curve

All controls are not always available in every plug-in processor.

For instance, only the L1-Ultramaximizer (shown as "+L1" in most application menus) and the L2 hardware mastering limiter offer complete control of all three elements listed above. The small IDR plug-in is internally set to type 1 Normal dither and curve, and the only user control is the Quantize, to select the wordlength.

Some plug-in processors have internal dithering, such as the Renaissance Equalizer and Compressor, which are dual-precision (48bit fixed on DSP, 64bit float on host), which is dithered using a simple flat dither for output to the system (dithered to 24bit/DSP, 32bit/host). No controls are available for this dither.

Other processors have internal noise-shaping (such as Q10, which has each band noise-shaped to provide 120dB signal-to-noise) which also has no user-available controls.

Chapter 3 - About Dither and Noise shaping

Dithering and Noise shaping are two independent, but complementary, techniques to improve the perceived quality of sound after it has been re-quantized.

As will be explained here in some detail, each technique is responsible for the improvement of a different subjective quality of the noise imposed by re-quantization. Therefore, each can be used separately to improve that specific quality.

Dithering is done in order to change the character of the quantization noise to more closely resemble analog hiss, rather than digital quantization noise. The main effect of dithering is to reduce (or, in case of type1, virtually eliminate) all correlation between the quantization noise and the original signal, thus reducing (eliminating) non-linear distortion typical of digital quantization noise. The dithering process 'exchanges' these distortions for a steadier analog-hiss quality signal.

Noise shaping is done in order to optimize the distribution of overall noise energy across the spectrum. This optimization is according to the ear's sensitivity. This means a decrease in noise (whether distortion or hiss) in the ear's sensitive areas (1 to 6kHz), is 'exchanged' for an increase of noise in less sensitive areas (above 15kHz, toward Nyquist).

Hopefully this has helped you see that in both techniques, the issue is about 'exchanging' the character and frequency content of noise (hiss & distortion) according to subjective criteria.

How do these processes help 'capture' 3 more bits of detail? The easiest analogy is to point to dithering in graphics, which is exactly the same process, and exactly the same type of psycho-perceptual model. The brain is capable of perceiving detail that is lower than a noise floor (in this case, dither). However, quantization noise is highly correlated to the signal; in other words, it is related to and governed by the signal. Dithering makes this noise become uncorrelated (as dither is a random signal), therefore allowing the brain to perceive the detail. The noise-shaping then helps to shift the energy of the noise to a less sensitive area of hearing.

Dither

1. No dither <off>. Within all applications (programs that Waves is compatible with and "plug into"), this setting means that the internal data is not dithered at all. In the case of the L2 hardware, it means that the internal 48bit data will be truncated to a 24bit output.The L2's IDR is more advanced, so technical information in the following paragraphs do not apply; refer to the L2 Manual for precise information.
2. IDR type1 dither. This dither is a wide-band dither. type1 adds a certain amount of noise causing a 5dB increase in background noise compared to no dither. It completely eliminates all low-level distortion and signal-dependent modulation effects. The result is a very transparent and clean low-level sound with a high resolution, most resembling the steady low-level hiss of an excellent quality analog system, and in lieu of digital quantization noise. This is the "purist" technology. It is designed for no nonlinear

distortion or modulation noise at low levels, and combines optimal dither noise with psychoacoustic noise shaping. type1 is also optimized to cause minimal side effects when used with stereo signals. type1 is the recommended choice for use when processing high quality mastering applications. By combining level maximization (peak controlling) and IDR processing, 16bit audio created from 20 or 24bit masters this way can have an apparent resolution of 19 bits, more than an 18dB improvement! When signals might be subjected to more than one stage of processing and quantization back to 16bits, the design of resolution-enhancement must satisfy additional requirements than a technology designed just for one-stage use. If applied several times in succession, a digital resolution enhancement technology optimized for one-stage CD mastering can produce unwanted side-effects. Waves type1 technology, however, is the first optimized for use at every processing stage, allowing for the effects of cascading and subsequent signal processing, if needed.

3. IDR type2 dither. This dither is a narrow-band dither, adds virtually no audible noise, and so is nearly 5dB quieter than type1, but with some low-level distortion. However, this distortion is generally much lower than with no dither at all. type2 is of a unique kind designed to minimize the amount of noise added, thereby giving a lower noise level than the IDR type1 process, but at the expense of some low level distortion. type2 does have some advantages for high quality mastering as well, and it is purely your choice whether the ultimate in low distortion of type1 is preferred, or the additional reduction in noise of type2.

Noise shaping

Another way to decrease the perceived level of noise is to "shape" the frequency content of the noise so it matches the ear's sensitivity curves. In basic terms, noise shaping shifts the noise to the frequency ranges where we hear it the least. The three options of Noise shaping provided in the +L1 and the L2 push more of the noise energy to higher frequencies above 15kHz and toward Nyquist, where our ears are least sensitive, and reducing the noise energy at lower frequencies. (Remember not all plug-in processors offer control of the noise-shaping curve). The three Noise shaping options progressively differ in the amount of this "shifting action".

1. Off. No noise shaping, resulting in more audible noise, (and distortion if dither is not used). The result has equal noise (distortion) levels at all frequencies, which is not optimal from a psychoacoustic point of view.
2. M (Moderate).
3. N (Normal). In addition to being very suitable for creation of Production Masters, using Normal with type1 dithering was also designed to be excellent for masters that would be processed again for any reason, including consecutive re-dithering, with an accumulation characteristic that is optimized to be minimal.
4. U (Ultra). This gives the greatest perceived hiss/distortion reduction. Ultra is a very high-quality setting, suitable only for use at the very last stage of mastering high-resolution audio (16bit and longer wordlengths) targeted for high-quality digital media. It is best to use Ultra in the last stage of audio preparation (Production Master). Due to the HF gain, it is theoretically possible that the relatively high amount of high frequency energy could cause undesirable side effects if the signal is going to be processed or digitally edited again. However, with many thousands of IDR-processed masters, no such situations have been reported or observed. Since it is theoretically possible, we wish to inform you of it.

You can hear the effect of Noise shaping by itself by monitoring the +L1 plug-in output and setting Dither type to Off, then selecting one of the Noise shaping options. Of course, the effect of Noise shaping is even greater when used with type1 or type2 dithering, since Noise shaping reduces the audibility of the added dither noise.

You may try the full effect of IDR technology by listening to the same material, with both IDR types and different kinds of noise shaping. The most obvious places to examine are notes or reverb during the end of the sound, or "tail". It is during this time that quantization error is most audible, although it is present on all low-level signals (such as elements that are soft in a mix, etc.)

If you don't feel you fully understand the tradeoffs between IDR and Noise shaping settings, the option that will generally work well for CD-mastering is type1 with Normal noise shaping. For minimum noise with 16bit and greater sources, type2 Ultra; for maximum resolution, type1 Ultra.

The noise reduction figures are even better for doubled sampling rates (such as those in the L2 hardware processor's IDR section). If audibility of noise were the only factor, the choice would almost always be to use Ultra noise shaping, but in some situations, heavy noise shaping of the Ultra kind can theoretically have some disadvantages, and the alternate settings, such as Normal or Moderate may be better.

One additional point that is not concerning a technical disadvantage of Ultra, but an aesthetic one, is that the Ultra setting can preserve and actually reveal the true quality of a master. Therefore, if the original master is not of great quality, you may prefer to use Moderate or Normal to better emulate other "less-revealing" wordlength reduction technologies, such as those that "smooth out" poor-quality masters.

For 16bit applications, Ultra shaping should be avoided in the following situations:

(a) Subsequent digital editing, when the signal is subjected to later editing. At the edit points, an extreme noise shaping might cause low-level, yet audible "clicks" in rare cases when played on very inexpensive CD players. An example application in which you would avoid use of Ultra shaping would be on CD's with production music or sound effects libraries that would certainly be subject to further digital editing.
(b) Poor Error Correction, i.e. when a signal is destined for a carrier medium with poor error correction, such as CDs pressed in pressing plants with poor quality control (who wants that anyway?). When errors that are not properly corrected occur, the Ultra setting, like all forms of heavy noise shaping and other resolution enhancement technologies, tends to cause audible background crackles, especially on very cheap CD players. While these effects generally don't occur on the majority of mid- or hi-fi CD players, they can be noticeable on very cheap products. The amount of such crackles on poor pressings is greatly reduced by the Normal noise shaping. Of course it can be argued that listeners with very cheap players, or users who don't care about poor quality pressings, may not be interested in high-quality reproduction at all!
(c) Subsequent production equalization before duplication (not by the user!), when heavy treble boost

equalization is subsequently employed. (Again, we're referring to EQ by a pre-mastering person before duplication, and not referring to a user who boosts the treble on their system. For the user to do so is fine, and is up to them.) This can cause the strongly boosted higher frequencies used by Ultra noise shaping to become so high in level that they might feed excessive noise energy into loudspeakers. Therefore Ultra shaping is best avoided in situations where subsequent equalization may be used in pre-mastering, such as in a compilation. However, if the mastering is done correctly the first time, large treble boosts would be quite unnecessary, and is somewhat moot in this context. Much less HF gain is used with the Normal and Moderate noise shapers (all of which could be observed on a real-time analyzer, such as Waves PAZ Psychoacoustic Analyzer).

Chapter 4 - Recommended IDR settings

Of course these settings are for when you have access to the full IDR control set, at present just the +L1-Ultramaximizer, or the L2 hardware.

Remember: The small IDR plug-in is internally set to type1 Normal.

Any combination of dither and noise shaping can be used, but the following settings are particularly recommended for different applications.

- General Purpose high-quality use, including material liable to be edited, EQ'd, and re-dithered: type1 Normal
- Lowest Noise: type2 - Ultra
- Low Noise/Highest quality (final production masters): type1 - Ultra
- Low noise while allowing editing/EQ: type2 - Normal
- High Quality, with lowest risks of spurious noises on edits or cheap CD players: type1 - Moderate
- Low noise, with lowest risk of spurious noises on edits or cheap CD players: type2 - Moderate

Read Me First

Installation
Tech Support

WaveSystem

L1
UltraMaximizer

C1
Parametric Compander

C1
Setup Guide

S1
Stereo Imager

Q10
Paragraphic Equalizer

Q10
Setup Guide

TrueVerb
Room Emulator and Reverb

PAZ
Psychoacoustic Analyzer

DeEsser

SuperTap
2-tap

IDR
(Increased Digital Resolution)

(North & South America)
306 West Depot, Suite 100
Knoxville, TN 37917
USA
Tel: 1-865-546-6115
Fax: 1-865-546-8445

(Rest of the World)
Azrieli Center, Tel Aviv
Israel 67011
Tel: 972-3-6081648
Fax: 972-3-6081656

WAVES

www.waves.com